InterActivity

TOOLS + TECHNIQUES FOR INTERACTIVE MEDIA DEVELOPERS

™

HOT WIRED

Building the coolest WEB SITE in the world

Panic in the Park

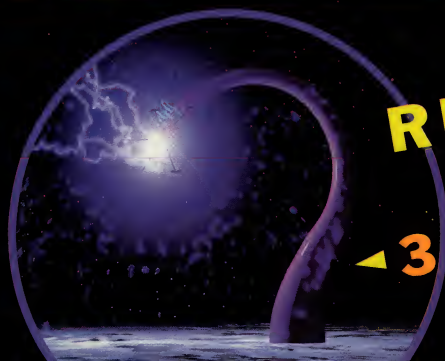A Production Powerhouse Reveals the Secrets of How to Produce Adventure Games in Six Months

REVIEWED!
Radius Telecast
◄3 ElectricImage Plug-Ins
Movie Cleaner Pro

InterActivity ™

FEATURES

HOW-TO COLUMNS

DEPARTMENTS

FROM THE EDITOR

TROJAN HORSE

Cue the cheesy theme song. Roll the flying logo. Pump up the hype and hysteria machine. The Headline News follows.

Talking head: "The Senate today passed sweeping reforms of the telecommunications industry. President Clinton has vowed to sign the Telecommunications Act into law. Some hail the bill as long-overdue legislation that will open the way for phone and cable companies to compete with each other."

Cut to pundit. "This means all the beautiful things we have been predicting would result from industry deregulation are going to come true. TVs as telephones. The phone company as a cable TV company. Cable as an Internet access medium. Home shopping online and on TV simultaneously via the wonders of ultra high-speed cable modems. More profits for all. And the consumer will be the ultimate winner because with increased competition will come lower rates. . . ."

Talking head. "Critics disagree."

Cut to pundit two: "This means all the beautiful things we have been predicting would result from industry deregulation are going to come true. TVs as telephones. The phone company as a cable TV company. Cable as an Internet access medium. Home shopping online and on TV simultaneously via the wonders of ultra high-speed cable modems. More profits for all. And the consumer will be the ultimate loser because monopolies are bad, but unregulated monopolies are worse. . . ."

Cue talking head. "Others say the big loser in the bill is pornography."

Cut to staff pundit. "That's true, Talking Head. The bill outlaws indecent materials, so pornographers beware."

Run video clip. "In other news, pop singer Michael Jackson was seen eating today with an unidentified woman just days after it was announced . . ."

To many of you this stuff is old news. Last issue we devoted a lot of space to the Telecommunications Act and its Exon/Coats Communication Decency Act (CDA) amendment and offered suggestions about what you could do about it.

But I still run into people who haven't got a clue that the CDA is something they need to worry about — people who see broadcasts on TV like the one described above and think, "Oh. They're regulating porn. That's got nothing to do with my business."

But there's some chilling stuff in the bill as it now stands. The CDA is a Trojan horse. Hidden in its anti-kiddy porn rhetoric are provisions that make anyone using television's Seven Deadly Words in an online forum that could be accessed by minors a felon.

More frightening, the bill makes criminal any content accessible to a minor that is "indecent or offensive or violent . . . as determined by contemporary community standards." Thinking about distributing your new *Doom* module online? Better plan on spending "not more than two years" in jail if some kid whose God-fearing parents think it's too violent or somehow satanic turn you in.

If you're operating a site that offers medical information, beware. If you include explicit information on how to avoid contracting HIV or talk about the horrors of abortion and your site is freely accessible to minors, once President Clinton signs this bill into law, you too can add felon to your resumé.

Of course, the list of organizations lining up to challenge the constitutionality of the CDA is long and impressive. But given how quickly this Trojan horse made its way through our elected officials, it's not safe to assume that the courts will shoot it down.

The best thing you can do is vote the bozos who drafted this law and all those who blindly supported it out of office next year, when nearly all of them are up for re-election. It's time we had elected officials whose understanding of technology didn't end with the steam engine.

End of tirade. Better go read our Case Study on how HotWired, the coolest Web site in the world, is put together before you need to send us your birth certificate to prove you're old enough to see it.

Dominic Milano
Editor

As your drummer takes the standard 92-minute solo, you wake up suddenly with this really cool idea.

The show ends. You head backstage, yell at the sound guy, and sit down at a Mac.

Look at you, you're an editor. Cut. Paste. Move this. Slide that. It's so simple, Sherlock.

Check out the video you did. Time to celebrate. Back to the hotel room. Hello, Security?

Rock and roll is here to stay.
But that prehistoric way of making music videos is history.
Hello, Media 100.® The most powerful and simple digital video system ever. Say goodbye to pricey studios and pushy editors.
Just plant yourself in front of your Macintosh® and go. So it's that simple. And oh-so powerful.
You get the highest picture quality available on any nonlinear system. Period.
Edit online, in real time.
You also get 8 tracks of real-time audio, graphics, titles, 100% QuickTime™ compatibility, stunning motion effects, and real-time ColorFX.™

MEDIA 100®

You got something in mind? Get it done—all on one system.
For a free video or live demo call 1-800-832-8188.
And instead of trashing hotel rooms, start trashing that old way of doing videos.

READER SERVICE NO. 3

W H A T   D I D   Y O U   H A V E   I N   M I N D ?™

P. O. V.

Web Checker

The following URLs listed in the article "3D Models and Texture Libraries" (January '96, p. 36) were not found: 3DSite, Gweb, and Mesh Mart. Have these been checked?

Bob Krawczyk
via the Net

The URL for 3DSite should have been printed as http://www.3dsite.com/3dsite/. We left off the final slash. The Gweb URL has changed. The new URL is http://www2.cinenet.net/GWEB/. We tried the Mesh Mart URL as printed and reached the site successfully.

Mad About Macs

Please cancel my subscription of *InterActivity*. I received my first copy (November/December '95) and read it today. The first article I read was "Interactive Storytelling, Part One" by Connor Freff Cochran.

Although his analysis of interactive games was fine, one statement was offensively inaccurate. He said the hell he described of his problems getting his PC hardware to run the CDs would be the same on the Mac: "Macintosh mavens should hold their tongues. Such headaches are just as common in the Mac world, different only in specific kind and not in general annoyance."

I follow the game market avidly and own most CD games available for the Mac. Out of 20 Mac CD games I have, the only problem I have ever had was a warning that the game would run faster if I had more memory.

The misinformation spread by uninformed, ignorant, PC -user writers is offensive, and I am increasingly less tolerant of it.

Cochran lost all credibility with this one statement, so much so that I found myself unable to finish reading his story. His article along with your magazine went straight to the trash bin (after finding your email address). I don't have the time to read articles written by people who won't check their facts before they publish them.

Joe Barrus
Vice President
IPD
via the Net

Any computer can make your life easier or more difficult. Personally I hate Macs because the OS interface is too simplistic. Drag and drop/click and hold/icon file management. Not for me. I like to see the filename extensions so I know what they are — i.e., *.txt, *.dll, *.exe.

Do you know how hard it is to get into the system files of a Mac? I think I spent 15 min-utes trying to change a memory configuration. As for CD-ROM programs not working, give me 20 CD-ROMs. I know they will install and run with no difficulty on my system (a 66MHz 486 with 16MB RAM and various peripherals running DOS 6.2a and Windows for Workgroups 3.11).

I can't remember the last time I had a real problem. The secret is to have a well-managed, maintained OS. Memory configuration is very important, and don't forget hard disk maintenance. If only manufacturers would ship a properly configured system. Yeah right. Let's all dream of plug and play.

Chris Marx
via the Net

Stock Options

Lea Anne Bantsari's excellent survey of stock footage suffered one glaring omission: Fox Movietone News. The Movietone archive, generally acknowledged to be the largest and most comprehensive of all the newsreels, spans the years 1919 to 1963 and encompasses some 55 million feet of film.

The archive is currently being transferred to videotape and digital data storage, and the card catalog is being converted to a CD-ROM based database. During the conversion, the archive is closed to researchers and potential licensees. The conversion project is targeted for completion in late 1996. Meanwhile, watch for Movietone News Online, which will debut shortly.

Dan Cooper
Managing Editor
Fox News
via the Net

Starry Alternative

The article about StarSight (January '96) fails to mention the precedent-setting scrolling onscreen TV guide by SuperGuide. I have enjoyed the benefits from this publisher of *OnSat* magazine for several years in my Uniden 4800 Super Satellite Receiver. I am aware that Uniden is offering StarSight in their latest satellite receivers; however, I saw nothing that new or innovative in your article. My UST4800 Super has an attached mouse that automatically operates my VCR to record easily selected programs. The SuperGuide datastream is downloaded from the satellite transponder for WGN.

• SuperGuide, Box 167, Shelby, NC 28151-0167; vox 800.234.0139 (subscriptions and renewal), 704.482.0491 (technical help); fax 704.484.7334; email superguide@shelby.net.
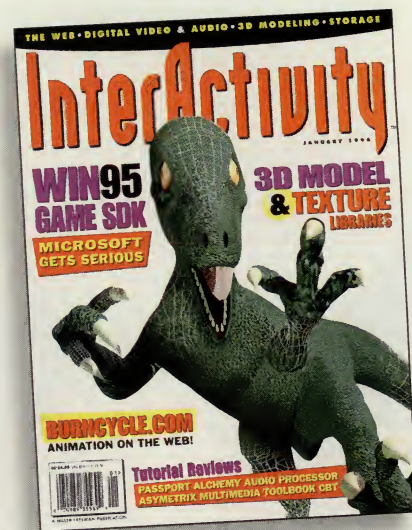
• Uniden America, 4700 Amon Carter Blvd., Fort Worth, TX 76155; vox 800.261.9498, 817.858.3300; fax 817.858.3849.
Raymond W. Polone
via the Net

School Daze, Part 2

I noticed that your FAQ column (January '96) regarding educational courses in multimedia did not mention a program at the Georgia Institute of Technology. Since 1991, the Continuing Education department has had certificate programs offering a range of courses in multimedia and Internet studies. Further information is available via our homepage: http://www.newmedia-coned.gatech.edu/.

Georgia Tech also has a master's program (M.S. in Information Design and Technology) offered by the School of Literature, Communication, and Culture.

Thomas Winn
Asst. Director of LCC ConEd
via the Net

Thanks very much for including UCLA Extension in your FAQ column (January '96).

I know the people responsible for the two other programs you mention in California — SFSU and CSU Long Beach. We all share the conviction that using professional practitioners to teach new media is essential. In our case, that includes *InterActivity* columnist David Javelosa, plus leading experts in authoring, editing, digital image creation, and so on.

Like SFSU, we offer a certificate program; our program is in film, television, video, and new media — level I and II. The program has two hands-on labs: the UCLA Extension/IBM Media Lab, based at Universal CityWalk in our metropolitan center (which will receive an upgrade to Pentiums with 64MB of RAM in the winter),

and the UCLA/UCLA Extension Music Media Lab, a Macintosh electronic music facility.

This past summer we initiated a new program in digital image creation, focusing on software that runs on the SGI platform. We're offering intro classes plus 14 workshops in winter '96, taught at the new Silicon Studio Training Center in Santa Monica. These classes, presented in cooperation with Weynand Training International, will be applicable to a new Certificate in Digital Image Creation, which is in the planning stages.

Additional courses in themed entertainment, virtual reality, game design, and other areas make it a comprehensive curriculum.

In addition to the vox number you printed, we can be reached at: fax 310.206.7435; Internet espa@unex.ucla.edu.

My congratulations to Dominic Milano and the entire *InterActivity* staff for an excellent new pub. Keep it up!

Charles S. Swartz
Program Manager
Department of Entertainment Studies and Performing Arts
UCLA Extension
via the Net

Kudos

Just a letter of praise for Chris Meyer's contributions to *InterActivity* (November/December '95). This is a real life example of how a good article keeps people reading a particular magazine. His review of CyberSound FX was excellent — right on target. I also found his "Creativity Under the Gun" one of the most informative articles I have read in a long time.

I just received the January issue (with my renewal card) and was happy to see that Chris was contributing again.

Keep up the good work and remember the Windows contingent.

David O. Hicks
Vice President
GeoGraphix
via the Net

Just want to say how impressed I am with the range of information provided by your magazine — especially "How to Put Together an Internet Web Site." I would love to get my hands on the first article in the series. Guess I will have to order the back issue.

Wish you had defined baud a bit more specifically — but everything else was clear as a bell. Enjoyed the Clement Mok piece too (November/December '95), though he got a bit too erudite to be understood.

I am a fan in Fresno.
Sharon
via the Net   ✺

RUN TIME

Compiled by Erica Smith

Better late than never. **Sharp Electronics** has decided to push efforts in the multimedia field and has formed the **Multimedia Development Center** (MDC). MDC will work on new product development and do research with the new Sharp Laboratories of America. Both will focus on virtual reality, digital video, and multimedia personal communications. . . .**Edelson Technology Partners** has decided to make a substantial investment in **The Duck Corp.**, original developer of TrueMotion video compression software and Comprending technology for videos and video games. . . .By the way, **SGI** has been taking some heat lately from **Intel** and **Intergraph** who crow that the new Pentium Pro processor is faster than SGI's entry level MIPS chip. Now Intel claims a bug in the software used to generate industry benchmark tests led to a 10% overstatement in results for the new Pentium and Pentium Pro processors in the SpecInt92 results. Some bug, huh? The problem was found as all the other chip makers tried to duplicate Intel's numbers and came up with different results. . . .*24 Hours in Cyberspace* is the newest title dreamed up by Rick Smolan, whose last work, *Passage to Vietnam*, was such a success. The high-concept formula from his *Day in the Life* series is now being applied to the Net. Photojournalists from around the globe have been chosen to submit work on how the online revolution is changing people's lives. A Web site will be up in mid-March to display all the participants' efforts. A CD-ROM, book, and TV documentary will follow. . . .Hooray for Hollywood. **Creative Arts Agency** (CAA) is getting hip. Teaming with ad agency Fattal & Collins, they're producing *American Cybercast*, an online network with dramas and sitcoms specifically tailored to the Net. Apparently three shows are in development, one being *The Spot*. CAA believes they'll be competing with primetime television, and, they hope, stealing some of the advertising dollars. . . .**3DO** is mov-

PYRAMID PLAY

What began as a programming exercise for Andrew Looney ended up as the strategy game *Icebreaker* (developed by Magnet Interactive Studios and distributed by Twentieth Century Fox) that recently hit the market. This is a game that wasn't created to fill a niche or make piles of money. It evolved out of his work because he was enjoying it.

Says Looney, "I was developing some software for a prototype of a game engine, and in doing so created a simple little test game. As soon as my co-workers got a chance to play it, they became addicted to it, and the next thing I knew, everyone was waiting to play."

The object of the game is to destroy all pyramids, and the game has 150 levels of play. No score is kept. Instead, you win by clearing levels. *Icebreaker* is more abstract than the usual "kill the enemy" games; with pyramids as the target, it's almost politically correct.

Originally done for the 3DO platform, *Icebreaker* was written in C++ using 3DO's

Icebreaker

software development environment, running under MPW on the Mac. The software was then ported to PC using Microsoft Visual C++ and to the Mac using CodeWarrior 5. The creative team used 3D Studio, Animator Pro, 3DO Animator, DeBabelizer, and Adobe Photoshop for the artwork and animation.

Looney chose to concentrate on simplicity rather than distract the player with features. One of his biggest challenges was creating the artificial intelligence for Seekers, moving pyramids that have to be gathered into a herd to be cleared and destroyed. Seeker pyramids are smart enough to maneuver out of the way.

Interactive WORLDS

WorldsAway is a collaborative online service from Fujitsu and CompuServe designed to create what promises to look like a shopping mall in Toontown. The main attraction is interactive chat rooms where users can take on an alternative identity and interact in visual environments. Users can

change their personas, create rooms, and furnish them. They can also buy real-world items and furnish their real-world environment as well.
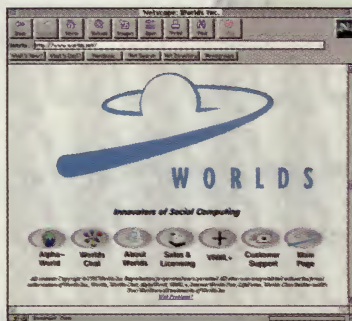
Early demonstrations of the software featured 3D rooms that look like Egyptian temples and avatars appearing in the forms of slow-moving creatures with cat heads, human heads, and so on. The implementation has its charm, but tedium seems to set in fairly quickly. The level of interaction is dependent on the sophistication of the users and, judging from most chat rooms, the prospect is not all that optimistic.

Some of WorldsAway's thunder has been stolen on the Internet by the AlphaWorld project from Worlds Inc., which was accessible before WorldsAway. The most notable locations in AlphaWorld are an art museum and a newsstand. Other possible structures include storefronts and gardens. Both WorldsAway and AlphaWorld plan to offer vendors the ability to create virtual stores. To visit WorldsAway, Go WorldsAway on CompuServe. AlphaWorld can be reached at http://www.worlds.net.

–Kathleen Maher

The HI-TECH GOLDRUSH Continues

ing to MPEG. Actually, they've announced a new product line for digital video compression. MPEGXpress, a Macintosh-based MPEG-1 encoding solution for developing content for a variety of platforms, is the first product to roll out. . . .Not the first, not the last. **Atari** joins a long but distinguished list of companies to form an interactive division. It'll be called — no surprises here — **Atari Interactive**. Their platform strategy calls for PC development first and foremost but also includes Saturn, PlayStation, and Mac. . . .The bedfellows just don't get any stranger than they are in Hollywood. Insiders are abuzz with the news that James Cameron's **Digital Domain** moved into the same complex as the **DreamWorks** studio in Playa Vista. As everyone knows, DreamWorks principal Steven Spielberg has a longstanding special relationship with Marin-based **ILM** headed by old friend George Lucas. Spielberg himself helped whip up rumors at a press briefing when asked about his loyalty to ILM. Spielberg answered, "I will stay loyal as long as ILM stays competitive." Most recently ILM has completed effects for Spielberg's **Amblin Studios** picture *Twister* and last year's *Casper*. Cameron, whose Digital Domain created the effects for *Apollo 13*, was careful to avoid all exposed toes but agreed that proximity held promise for his growing effects house. . . .The founder of **ESPN**, Bill Rasmussen, is starting up the **Attitude Network** with **Internet America** president David Rae. Comparing it to cable television, Rasmussen pledges to deliver customized programming for niche lifestyles. Their first program is **Happy Puppy**, a popular games site they have acquired. . . .Sad news for many colleagues. **Digital Pictures** laid off 30% to 40% of their employees

Companies are going public at an alarming rate. At least it's alarming in the high-tech industry where companies making IPOs are taking off like rockets — to the consternation of those who've followed these companies for years. According to the San Jose *Mercury News*, 68 companies in Silicon Valley went public in 1995. That's more than twice the 1994 total of 31 and breaks the previous record of 57 nationwide.

But as the stock market went crazy for high tech, reality set in. SGI's earnings were lower than expected, which CEO Ed McCrackin blamed on a slowdown in sales to the U.S. government and lower revenues from OEM channels. Adobe took a hit on their acquisition of Frame Technology and admitted that nonetheless revenues were lower than expected. High flying Micron has taken a 60% dip. And, of course, Apple's shortfall has been well documented. But these examples aren't the only indicators of a slowdown. New Hampshire-based IDC predicts that the shakeout of online and Internet related businesses will continue for most of this year, and Dataquest has documented slower sales of home multimedia PCs over the last quarter of '95 and the first quarter of '96. And here in Silicon Valley, it's hard to miss the decreasingly generous multimedia gravy train as old friends do the job hop.

So how to explain the trend of IPOs? In spite of the handwringing, the overall strength of high-tech stocks has propelled the IPOs into the big time. According to the San Jose *Mercury News*, the average IPO rose 37%, while the average Silicon Valley IPO rose 64%.

Pixar gambled big time with *Toy Story*, spending more than three years in production and using their own rendering technology to do it. It worked. Pixar's revenues shot up 145% as a result of the film's success, and their initial stock offering was one of the storybook tales of the 1995 stock market. Before this however, Pixar went through some pretty rough times. Renderman, Pixar's rendering library capable of creating spectacular photorealistic images, had a brief heyday a few years ago. Their success met with trouble when software companies that OEMed the product had trouble selling it because it was too slow. Pixar smartly went into the production business themselves and have managed to demonstrate the power of both their software and their talent. The infusion of money from the IPO could help Pixar deliver on their early promise, but the odds are just as good that they could stall at the gate. They won't have three years to produce the next hit. The *Toy Story* merchandising machine will probably give the company breathing room for now.

Some see Netscape as another example of a darn scary IPO successful beyond all reason. Their IPO rose 396%, but the company wasn't selling their product — they were giving it away. On the plus side, Netscape is ubiquitous on the Internet and, to their credit, the Netscape executives have been plowing money back into research. They've been the first to incorporate VRML, RealAudio, and various third party plug-ins such as Macromedia's animation format to their Net browsers, and their authoring tools are winning them the hearts and minds of the growing army of Web masters. But what about that Internet slowdown predicted by Dataquest? Is Netscape going to be able to deliver to skittish investors?

The latest in bold IPOs is a little Massachusetts company called Spacetec that specializes in 3D motion control input hardware controllers and software systems. The company's 3D-I, a ball-shaped input device that lets users use the mouse to manipulate 3D graphical images in real time as if they were holding an object (or moving objects through space) is a big hit. This technology (along with that of Logitech's Magellan) has been endorsed by Alias/Wavefront, which is advocating the development of more intuitive two-handed manipulation of 3D objects in its Project Maya interface design. It's got plenty of promise, but it's still a pretty expensive input device.

Opinions are divided on the message sent out by the rampant public offerings. Some worry about the meteoric climb of the stocks. Others anticipate slower growth but nothing resembling collapse. We're in a period when the public is so enamored of technology that they'll go hog-wild for stock from unproven companies, yet they won't buy the thousands of new CD-ROMs being released and they're price sensitive when it comes to hardware.

–Kathleen Maher

Once *and* Future GAMING

after a disappointing fourth quarter last year. And they're not the only company going through or preparing for a downsizing. . . . Here come the DVD players. **Toshiba** is reportedly shipping its first two models in the fall priced at $600 and $700. **Thompson Consumer Electronics** has also announced they will ship their player at a price of $500 in September. **Columbia TriStar** promises to have 50 titles ready as soon as the players ship. . . .**NTT Software Corp.**, a Tokyo-based company, has tested their 3D virtual reality concept in the form of Interspace software, which combines video-conferencing with 3D graphics environments. NTT ran a pilot program called CyberCampus recently at San Francisco State University and the College of Marin. Once on CyberCampus, users can click on "teleportals" that move them through InterSpace. The options are a Levi Strauss site, Tower Records store, a tour of Marin County, and visitor information that describes NTT. By the beginning of April '96, NTT hoped to have 100 corporate-sponsored workstations up and running. . . .RAD Software has acquired the distribution rights to Miles Design's AIL sound library product line. AIL, commonly known as the Miles drivers, will be renamed the Miles Sound System and is available for DOS, Windows, and Win32. RAD expects to port to Mac soon. . . . For those who haven't looked at Netscape's homepage in a while, the folks who make the most popular browser in the world are hosting the first Java and JavaScript conference for professionals creating live online applications and content. It's scheduled for March 5-7 in San Francisco. 🐝

CyberCampus

Got News? Send it to *InterActivity*, 411 Borel Ave., Ste. 100, San Mateo, CA 94402; fax 415.655.4360; Internet interactivity@mfi.com.

While Sega and Sony battle it out head to head, Nintendo is sitting pretty in the back seat. You can read this two ways: Either they're biding their time until they're ready to unleash the Ultra 64 juggernaut, or the once-dominant market leader has lost its edge. Most pundits are taking the latter position because it's more fun, but certainly the others are keeping one eye on Nintendo while they duke it out. Sony launched the PlayStation in the U.S. in time for Christmas '95, claiming to have shipped 800,00 in North America by the end of the Christmas '95 season. At that time, Sega's Saturn was running in second aided by consumer awareness of its arcade titles, reporting sales of 400,000 in North America (3 million worldwide) by the end of 1995.

And what about Nintendo? They came clean about the delay of the Ultra 64 last spring at E3 (the Electronic Entertainment Exposition). Now they're assuring the market that when the new platform gets here,

Nintendo's Ultra 64

it'll be accompanied by a slew of games. Nintendo expects approximately 20 titles to be available by the end of 1996, and nearly half of those should be ready by the launch. For instance, Electronic Arts has announced an Ultra 64 version of its international best-seller *FIFA Soccer*. Expect to see plenty of these announcements before the platform's expected release this month.

According to Nintendo, the platform itself is just fine, and the technology codeveloped with SGI's MIPS division is ready to go. Now they're just waiting for the buildup of a critical mass of games worthy of such a hot new platform. Sega of America president Tom Kalinske put his own English on Nintendo's spin. In a recent issue of *Computer Entertainment News*, he questioned his rival's commitment to cartridge-based gaming. Kalinske suggests that Nintendo is buying time while they develop a CD-ROM Ultra 64 platform. Regardless of what Nintendo unveils in April, they've got a lot of catching up to do after missing Christmas '95. 🐝

–Kathleen Maher

The Sundance Film Festival has helped redefine American filmmaking. Can it do the same for multimedia? Aside from the festival itself, the Sundance Institute founded by Robert Redford provides training and seed money to independent filmmakers. For this year's festival in January, Apple teamed up with the Institute to create a New Media Center, which highlighted the potential of computer effects and multimedia with demonstrations of Enhanced CD, special effects techniques, VR, and more. The festival also moved to a new dimension courtesy of Apple's QuickTime Live, which brings video and animation online. The festival used QuickTime Live to present interviews with filmmakers at the festival.

DANCING WITH APPLE

Apple and Sundance have signed a three-year contract that establishes the New Media Center as an Apple resource. Of course, Apple products including QuickTime, QuickTime VR, and the Apple Internet Connection Kit will be featured, but the Center also will highlight products heavily in use in the entertainment industry, such as Premiere, After Effects, and Page Mill from Adobe; Macromedia's line; Strata's Studio Pro; and Storyboard Artist, Digital Box Office, and Costume Pro from Power Production Software.

As for the multimedia industry, a little new blood can't hurt. It appears that Apple is striving mightily to protect its niche in the entertainment development industry, while announcements of less than spectacular earnings and layoffs come fast and furious. 🐝

–Kathleen Maher

PUTTING TOGETHER AN INTERACTIVE MULTIMEDIA SYSTEM IS LIKE OPENING A CAN OF WORMS*

Unless you offer your client the single-source CD-i solution from I²m.

When your client demands totally integrated full-motion video, stereo sound, graphics, animation and narration all in one system, I²m's Media Playback expansion card is the answer. Media Playback turns PCs and Macs into professional quality CD-i players, permitting use of monitors, modems, disk drives, printers and mice.

MediaPlayback, because serious interactivity demands serious hardware.

I²m
INTERNATIONAL
INTERACTIVE
M E D I A

275 Wyman St., Waltham, MA 02154   tel: (617) 290-5964   fax: (617) 890-6682   internet address: I2M@World.std.com

N E W G E A R

Compiled by Ted Greenwald

VIRTUAL SOUNDCARD for Mac

In music production facilities throughout the world, the personal computer has evolved from scheduling and accounting aid to multitrack recording deck without missing a beat. Now, with software from InVision entitled CyberSound VS ($249), just about any Macintosh becomes a high-quality music synthesizer.

CyberSound VS, dubbed a "virtual MIDI soundcard" by the manufacturer, comprises 50MB of 16-bit audio data — 512 synthesizer sounds created using methods from subtractive synthesis to physical modeling — and provides MIDI control of between 24 and 128 voices (the latter on a Mac 9500). CyberSound VS also provides reverb and chorus plus the CyberPlayer 16-track sequencer, effectively turning an Apple PowerBook into a MIDI studio. The software supports the General MIDI (GM) and Roland GS sound sets; OMS, FreeMIDI, and MIDI Manager system extensions; Apple Sound Manager; and QuickTime. A free demo is available.

• InVision, 2445 Faber Pl., Ste. 102, Palo Alto, CA 94303-3316; vox 800.468.5530, 415.812.7380; fax 415.812.7386; Internet http://www.cybersound.com, dennym@cybersound.com.

Reader Service #110
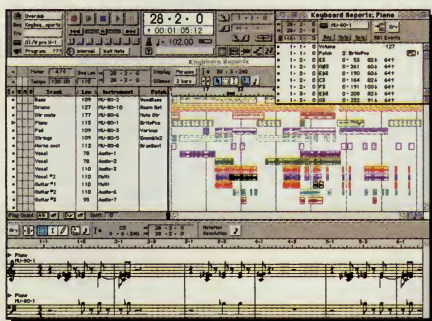
Audio/MIDI INTEGRATION

When it was announced at the 1995 NAMM show, Opcode's Studio Vision 3.0 ($995) for the Mac represented a new level of power in music production tools. Now you can get your hands on it. The new version of this pioneering MIDI/audio application provides unprecedented integration between the two disparate technologies thanks to functions for converting between audio and MIDI. For instance, audio pitch can be mapped to MIDI note number, amplitude to MIDI velocity, and event length to MIDI note-on and -off timing. In addition, the software automatically detects and corrects variations in the tempo of an audio track without introducing pitch artifacts. Conversely, you can translate an accelerando or ritard from a MIDI se-
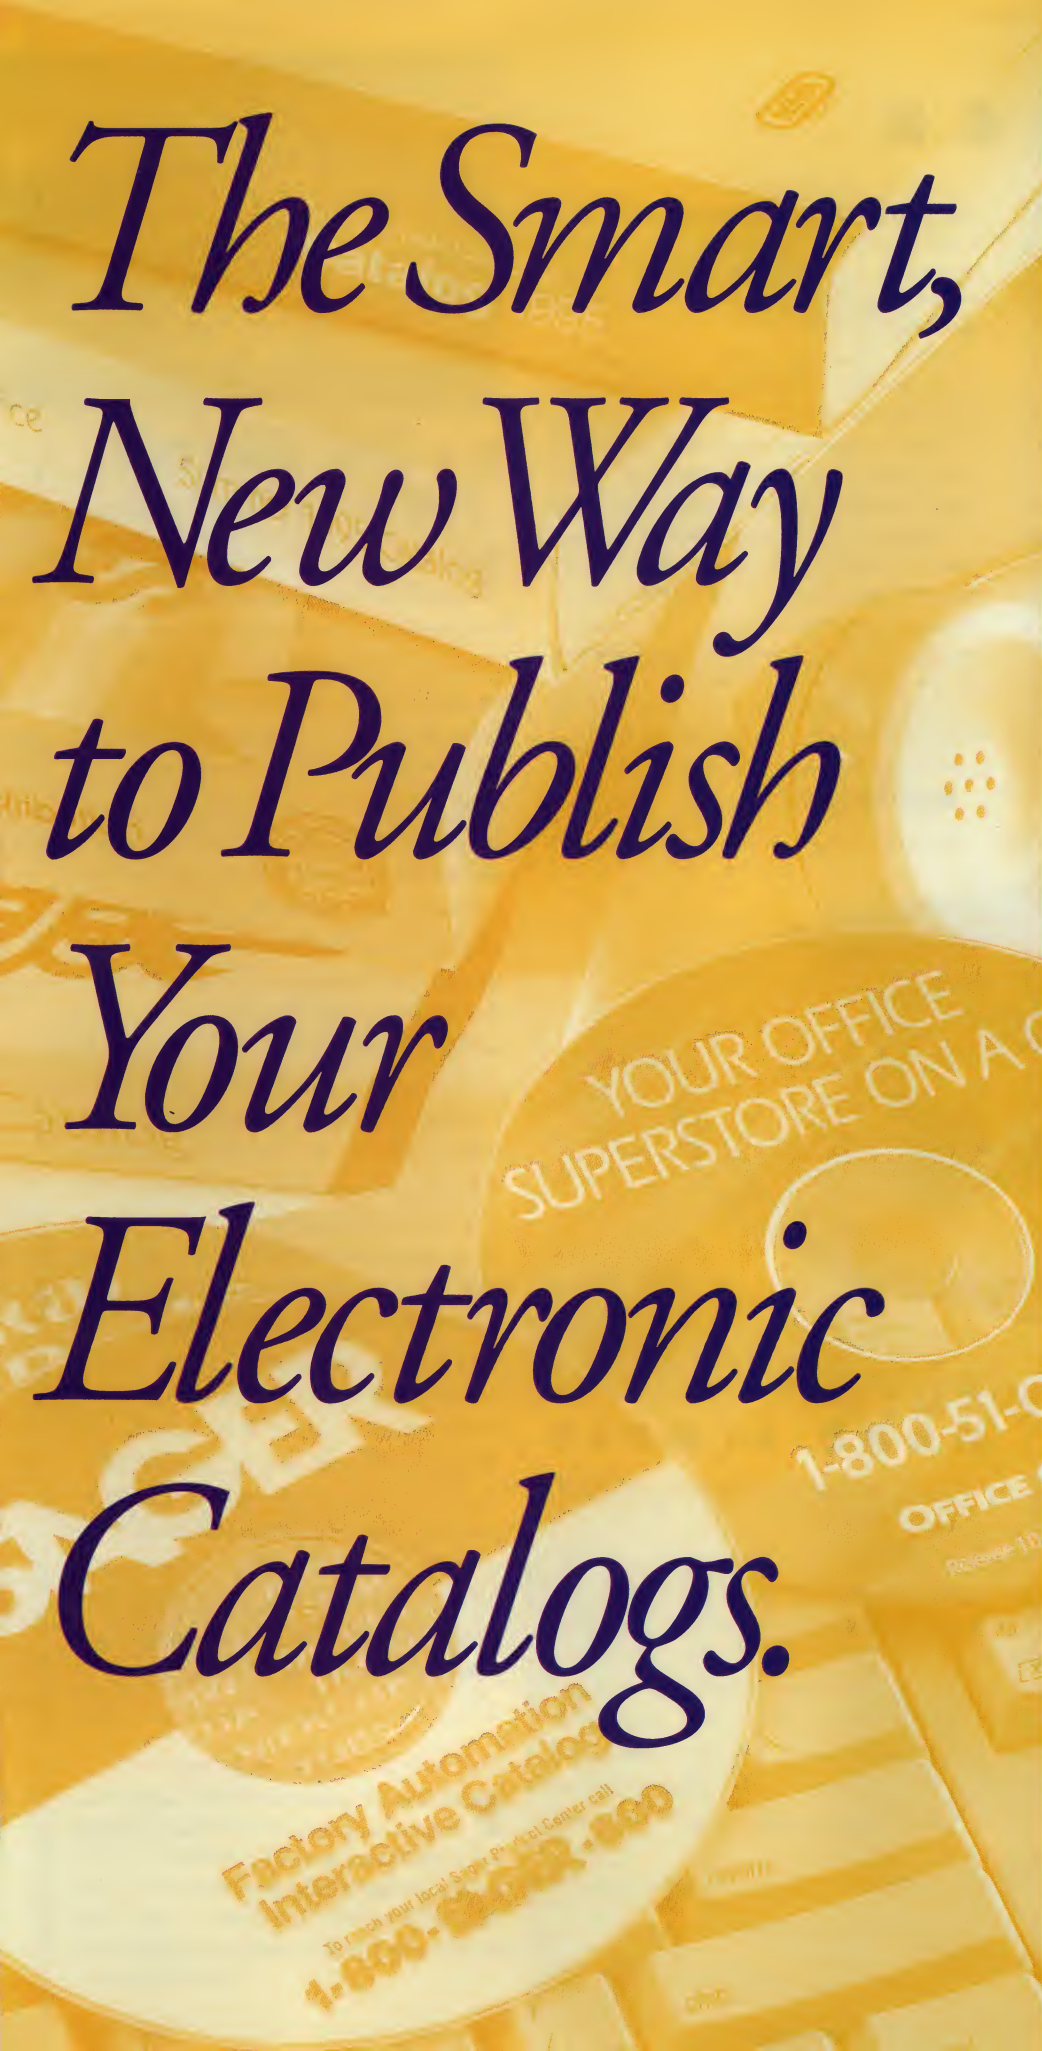
Export TO Saturn

For beleaguered game developers trying to accommodate multiple platforms, Nichimen Graphics offers time-saving software modules for N•World, their suite of SGI-based game development tools. Hot on the heels of Action Editor, which converts assets developed in N•World to Sony PlayStation, comes Saturn Express ($3,000), which converts object and texture data from N•World to Sega Saturn. The program supports vertex, polygon, and attribute data; Gouraud shading tables; and texture maps of any size; and enables platform-specific attributes to be assigned. Objects can be previewed as solids or wireframes on both platforms simultaneously, and camera angles set in N•World update instantaneously on the Saturn. Objects are exported directly in SGL or SGA3 format.

• Nichimen Graphics, 12555 W. Jefferson Blvd., Ste. 285, Los Angeles, CA 90066; vox 310.577.0500; fax 310.577.0577; Internet http://www.nichimen.com.

Reader Service #111

Continued from page 14

quence to an audio file without affecting pitch.

Studio Vision 3.0 is compatible with Digidesign Pro Tools III and the Digidesign TDM bus as well as Yamaha CBX-series hard disk recorders. It also supports Opcode's own DSP plug-ins for time compression/expansion, pitch shift, volume normalization, EQ, sample-rate conversion, and the like. Running on a Power Mac, two to 16 tracks of audio can play simultaneously without additional hardware.

New versions of other Opcode Mac programs have arrived as well. The Vision 3.0 MIDI sequencer ($495 bundled with Galaxy 2.0 universal synthesizer patch librarian and OMS 2.0 system extension) supports digital audio via Apple Sound Manager, enabling it to record, play, and edit MIDI and audio together running on a Power Mac without extra hardware. The new version also features enhanced integration with Galaxy and OMS. For example, when you load a sequence, Vision will determine whether your synthesizers hold the necessary patches; if not, Vision calls Galaxy to load them. (The update makes Studio Vision AV obsolete, so it has been discontinued. Owners of the AV version can upgrade to Vision 3.0 or Studio Vision Pro 3.0.)

OMS 2.0 (that's Open MIDI System) has been given a facelift to streamline studio setup, and MIDI interfaces are now detected automatically. In addition, a new QuickTime Music OMS driver gives sequencing applications access to the Mac's resident GM sound set, making it possible to audition sequences without being connected to external MIDI instruments. Galaxy 2.0 uses the new Name Manager facility of OMS to enable the names and numbers of patches in Galaxy bundles to be displayed within Vision, where they can be searched and accessed.

• Opcode, 3950 Haven Dr., Palo Alto, CA 94303; vox 415.856.3333; fax 415.856.3332; Internet http://www.opcode.com.

Reader Service #112

If a 3D Object
Fell in the Forest...

The latest version of Presenter Professional 3.1 ($1,995), a Mac-based 3D authoring system from Visual Information Development Inc. (VIDI), incorporates a set of advanced audio tools called Digital SoundStage. Designed to enable multimedia, game, and broadcast producers to simulate a real-world production environment, Presenter Pro offers a "3D" stereo field as well as naturalistic motion effects such as Doppler pitch shift. Virtual microphones can be placed anywhere in 3D space and assigned a sensitivity, range, directional pickup pattern, even an animated motion path. Sound elements can possess the attributes of distance, velocity, and direction. Thus, sounds associated with moving objects onscreen exhibit appropriate dynamic characteristics, making for enhanced realism and immersive depth.

In addition to these unique audio facilities, Presenter Pro offers a variety of multimedia capabilities. Video clips can be projected onto still and moving objects. Views generated by multiple virtual cameras can be animated and edited to create multicamera shots. 3D objects can be created using spline-based and solid modeling tools and given attributes including breakability, center of gravity, collision detection, elasticity, gravity, jiggle, mass, momentum, and wind. Furthermore, motion plug-ins can be applied to create effects such as flocking, tracking, and linking. High-speed ray tracing is included along with support for Pixar's RenderMan rendering application.

• VIDI, 136 W. Olive Ave., Monrovia, CA 91016; vox 818.358.3936; fax 818.358.4766; Internet vidisales@aol.com; AppleLink vidi.

Reader Service #114

Motion Capture
MADE SIMPLE

MotionStar, the most recent model in Ascension's Flock of Birds line of motion tracking systems, is designed for portability, ease of use, and robust operation. Performers move freely, their motions detected via an extended-range DC receiver. DC technology enables performers to move among obstacles, unlike video-based systems, and to move over a wide area while minimizing distortions common to AC-based systems. Captured motion can be applied directly to animations using off-the-shelf programs from Alias/Wavefront, Softimage, Autodesk, and others.

MotionStar, which comes in various configurations starting at $19,315, packs electronics, power supply, and a choice of interfaces (RS-232, RS-485, Ethernet, SCSI) within a single rackmount box. The enclosure holds up to 20 cards, each supporting a single sensor. Each sensor is sampled 144 times per second and multiple sensors are processed in parallel for fast tracking and efficient rendering within the host computer. Several MotionStars can be connected in order to track a total of 60 sensors simultaneously for interactions among multiple animated characters.

• Ascension, Box 527, Burlington, VT 05402; vox 802.860.6440; fax 802.860.6439; Internet ascension@world.std.com.

Reader Service #113

Win95 Disk Array

FWB continues to diversify its Hammer line of data storage products. The latest addition is not only the company's first Windows product but the first RAID 0 disk array for Windows: the PCI Fast & Wide SCSI-2 array for Win95 and Windows For Workgroups (price TBA). The same

technology for WinNT is expected by the time you read this. Optimized for Pentium-based computers, the new model comes in 2GB, 4.1GB, and 8.2GB configurations designed for high-throughput applications.

Also, a 2GB model has been added to FWB's mid-level line of hard disks, the hammer•PE 2000 available in internal ($989) and external ($1,079) versions. Average seek time is 8.5ms, while sustained transfer rate ranges between 4.45MB/sec and 8.0MB/sec (via Fast SCSI-2). Rugged construction enables the unit to be moved on a regular basis, and it's backed by a five-year manufacturer's warranty.

Meanwhile, FWB announces a new price for the hammer•CD-R 2x CD recorder, which now retails at $1,499. To facilitate backup and archiving, Dantz Retrospect 3.0 and Dantz CD-R Driver Kit for the Mac have been added to the bundle along with Toast CD-ROM Pro 2.5 and FWB CD-ROM Toolkit 2.0.

• FWB, 1555 Adams Dr., Menlo Park, CA 94025; vox 415.325.4392; fax 415.833.4657; Internet http://www.fwb.com, info@fwb.com; AppleLink fwb.

Reader Service #115

VIRUS Protection
For WinNT, Win95, Mac, and Unix

Dr. Solomon's Anti-Virus Toolkit by S&S, previously for Windows 3.1, DOS, NetWare, and OS/2, is now shipping in versions for Windows NT, Win95, Macintosh, and SCO Unix ($99-$125). The software detects and removes more than 7,000 computer viruses including the new Microsoft Word macro viruses. It works via WinGuard, a virtual device driver that transparently checks every disk and file accessed, including compressed and archived files, to prevent infected programs from running. A newly updated version of the driver that eradicates the AOLGold Trojan Horse can be downloaded from S&S's Web site.

S&S, 17 New England Executive Pk., Burlington, MA 01803; vox 800.701.9648, 617.273.7400; fax 617.273.7474; Internet http://www.drsolomon.com, info@ us.drsolomon.com; BBS 617.229.8804; CompuServe go drsolomon.

Reader Service #116
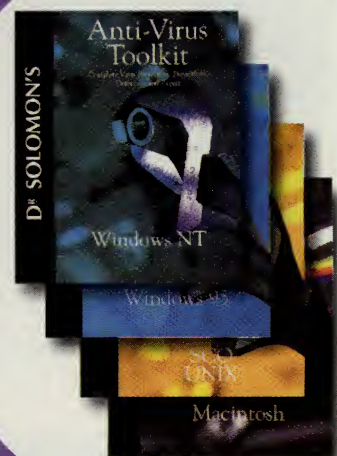
The Next Video Revolution?

Touted in some circles as the next Video Toaster, the Play Trinity ($5,995) is designed to put the functionality of a network-level broadcast video studio on the desktop. The hardware/software system takes advantage of the power of the Pentium processor, Win95, and PCI bus performance to deliver a 10-input, all digital, D1 video switcher at a dramatically reduced price point. The switcher accommodates eight live inputs and two still-store channels, each with its own independent color correction and effects, with a 16.8 million-color background matte generator and special effects galore.

The chromakeyer creates independent chroma or luma keys or a unique ChromaLuna key with simultaneous luminance and chrominance ranges. In conjunction with a programmable color effects engine, it's possible to remove all colors from a scene except a defined range. Other color effects include sepia tone, monochrome, day-for-night, solarization, embossing, posterization, and negative color.

The system includes two independent video effects channels with complementary sound effects. Live video can be processed at 60fps without chunky or blocky artifacts. For an additional $7,995, the Trinity Warp Engine provides 3D video effects such as mapping video onto 3D objects in real time with ripples, waves, swirls, cubes, crumpling, peeling, spinning, shattering, or pouring effects. Up to eight simultaneous 2D and 3D effects can be combined.

The character generator provides 10 nanosecond resolution, variable scroll/crawl speeds, realtime effects for keyed titles, and a host of high-end features. Paint, compositing, and animation facilities are also included with realtime airbrushing on shaded 3D surfaces, natural media brushes such as watercolor and chalk, anti-aliasing with adjustable alpha-channel transparency, and real-world physics including velocity, wind, gravity, and friction.

The timeline editor is designed for graphic drag-and-drop operation with an upgrade path to nonlinear editing via the Preditor two-channel hard disk option ($4,995). EDLs can be viewed in timeline, storyboard, or text formats, and 50 levels of undo are available for edits and effects. Automated audio mixing is included with EQ, effects send/return, and stereo panning. Eight balanced XLR or unbalanced RCA inputs are provided plus one stereo output (also XLR or RCA). Integrated machine control makes tape-based and disk-based sources interchangeable. Four RS-422 ports are provided along with direct support for common VTR models.

The Trinity system can be configured with inputs, effects channels, video formats, and so on to suit a variety of production needs. It boasts an open hardware architecture and object-oriented operating system, presenting an opportunity for third-party developers to provide plug-in software and other add-ons. Specs, source code, and documentation are available.

• Play, 2890 Kilgore Rd., Rancho Cordova, CA 95670-6133; vox 800.306.7529, 916.851.0800; fax 916.851.0801.

Reader Service #117

1000 Points of LIGHTWAVE

Light-ROM 3 ($49) is a triple CD-ROM set of 3D models, textures, and other graphics resources from Graphic Detail. It's available only until June 1, 1996, after which it will be replaced by a new set.

The highlight is thousands of objects and scenes for NewTek's LightWave modeling, animation, and rendering program, users of which are the target of these discs (but elements in various other formats are provided as well). For instance, 175MB are devoted to objects formatted for Impulse Imagine, 100MB for Autodesk 3D Studio, 30MB for Byte By Byte Sculpt, and 7MB for Realsoft Real 3D. Also included are 3D landscapes in LightWave, Imagine, and 3D Studio formats, plus 1,000 digital elevation maps (DEMs) for Virtual Reality Labs VistaPro, Natural Graphics Scenery Animator, and Questar World Construction Set. Not to mention 700 JPEG-format textures, a directory for the NewTek Video Toaster containing wipes and CG fonts, and a batch of useful public-domain programs and text files for Amiga and Windows platforms. Thumbnail renderings of all objects, textures, landscapes, and DEMs are provided for easy browsing.

• Graphic Detail, 4556 S. 3rd St., Louisville, KY 40213; vox/fax 502.363.2986; Internet michael@iglou.com.

Reader Service #118

Internal CD-R

CD-R technology is ready for the desktop with MicroNet's internal recordable CD drives for PCs, the 2x Master CD Plus Internal ($979) and 4x Master CD Pro Internal ($2,595). Based on a Ricoh mechanism, the Master CD Plus fits in the half-height 5 1/4" bay on any standard PC and provides a data transfer rate of 300kB/sec. The Master CD Pro, based on a Yamaha mechanism, boasts a data transfer rate of 600kB/sec and writes an entire CD in 19 minutes. Both drives come bundled with Corel CD Creator premastering software.

Meanwhile, the list prices of the equivalent external MicroNet models have been reduced. The Master CD Plus External drops from $1,895 to $1,295. The Master CD Pro External drops from $3,995 to $2,895.

• MicroNet, 80 Technology, Irvine, CA 92718; vox 714.453.6100; fax 714.453.6101; Internet http://www.micronet.com.

Reader Service #119
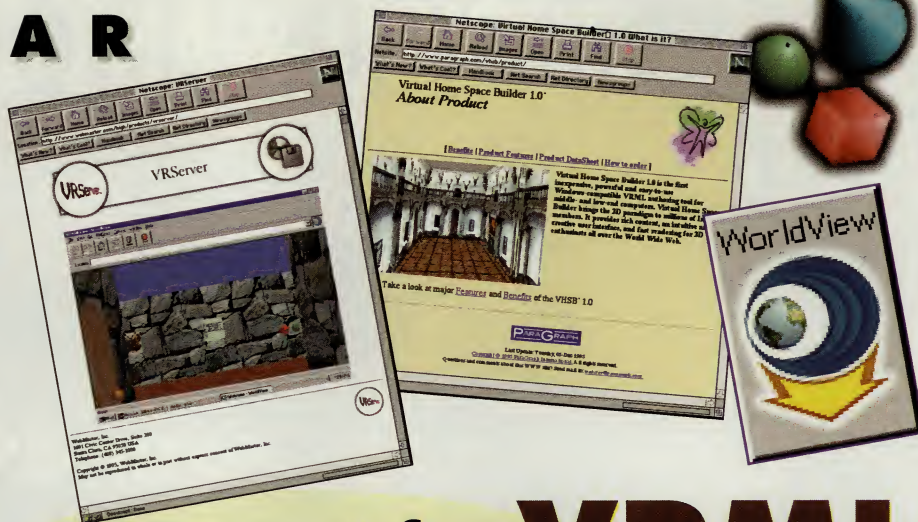
Support for VRML

You can almost see the bulges protruding as the World Wide Web, already stretched to the limit by an expanding base of users, imaginative developers, and shrewd marketing executives, begins to bust out in 3D. Thanks to strategic alliances among numerous movers and shakers in the interactive arena, software is proliferating that promises to make the Internet a truly multidimensional experience. In the short run (at least), the keys to 3D on the Web seem to be support for InterVista's VRML (Virtual Reality Modeling Language) file format, backward compatibility with HTML, ease of use, and compatibility with Windows.

InterVista itself is giving the technology a push with the first standalone VRML browser, WorldView ($49) for Windows or Mac. A pre-release Windows version is available for download as we go to press, enabling surfers to navigate VRML-enabled sites from a first-person perspective. WorldView also functions as an add-on to Netscape Navigator; when Navigator encounters a VRML environment, it automatically calls WorldView and exchanges data with it seamlessly. (Once you've installed WorldView, check out the first VRML site on the Web, Sony Pictures' promotion for the film *Money Train* at http://www.moneytrain. com/home/mtrain.)

InterVista has entered into alliances with the likes of Microsoft, Apple, Intel, Creative Labs, NewTek, and Caligari to foster acceptance of the format. One of their partners, ParaGraph, offers the first VRML authoring program, Virtual Home Space Builder (VHSB) 1.0 Commercial Edition ($495) for Windows. ParaGraph offers the final beta version for $49 and a time-limited final beta via their Web site free of charge.

Designed to accommodate nonprogrammers, VHSB creates 3D Web sites, or *home spaces*, that can be saved either as VRML files or in ParaGraph's proprietary MUS format. In addition to realtime interactive 3D graphics, sites can include text, audio (WAV and RealAudio), MIDI, video, and animation (AVI and FLI) — in fact, ParaGraph suggests that VHSB makes a good presentation authoring tool as well. Any part of a 3D environment can act as a URL link, and executable function modules can be attached. Home spaces can be decorated with textures, including tiled and animated textures, and with transparent colors. The software comes bundled with hundreds of world templates, textures, movies, and sound files, and custom content can be imported easily. A Virtual Home Space Viewer is included.

Another step toward the ubiquity of VRML-enabled 3D worlds on the Web is VRServer ($249) from WebMaster. Running under Windows NT, VRServer automatically converts HTML sites into VRML spaces, vastly simplifying the creation of 3D Web sites. The program provides control over lighting and placement of objects, and it ships with preconfigured spaces, called *themes*, including a medieval castle, Victorian mansion, and spaceship. The resulting space can be viewed using any VRML browser while the original HTML pages remain available to HTML-compatible browsers. Furthermore, the VRML environment is updated automatically when changes are made in associated HTML pages.

• InterVista, 303 Sacramento St., 2nd Fl., San Francisco, CA 94111; vox 415.434.8765; fax 800.317. 4204, 415.989.9910; Internet http://www.intervista.com, info@intervista.com. **Reader Service #120**

• ParaGraph, 1688 Dell Ave., Campbell, CA 95008; vox 800.810.0055, 408.364.7700; fax 408.374. 5466; Internet http://www.paragraph.com. **Reader Service #121**

• WebMaster, 1601 Civic Center Dr., Ste. 200, Santa Clara, CA 95050; vox 408.345.1800; fax 408.247. 9372; Internet http://www.webmaster. com, info@webmaster.com. **Reader Service #122**

CASE STUDY

Champagne Cocktail
drink of the week
The Alchemist
Suggested Servin
The Virtual

Manic Love

HOT

seinfe
Gallery
who is th
ines
utlan
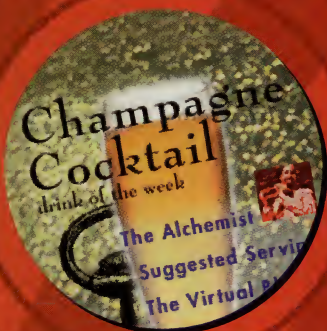
A BEHIND THE SCREENS LOOK AT HOW THE

BY MICHAEL MACRONE

In mid-1994, there really wasn't much of anything on the World Wide Web you could call a "magazine." Today there are dozens, if not hundreds — nobody, not even the indexers at Yahoo, knows for sure how many. What made the difference? ● To cite the Web's newfound popularity doesn't explain much, given that it was partly caused by the growth of online publishing. It's possible, however, to find a few more tangible factors in the Web's evolution from the beginning of '94 and beyond. ● One factor was the growing engagement of journalists with the Net. Not only were more articles written about the Internet, but more articles were being written for it in the form of bulletin board postings, Usenet articles, and Web pages. Writers were starting to see the Net as a medium as well as a subject. ● A second factor was the release of Mosaic Netscape, now Netscape Navigator, the dominant Web browser on the

CYBER RIGHTS NOW

WIRED
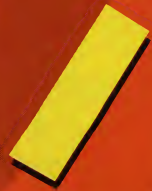
WORLD'S COOLEST WEB SITE IS PRODUCED

market. With Netscape, many users were finally able to see the point of the Web; the client was faster, having been optimized for slow connections, and it made pages look a lot nicer. It also offered authors new tools for controlling page design, principally through extensions to HTML, the encoding language of Web documents.

A third factor was the debut, on October 27, 1994, of HotWired.

■ ■ ■ ■ ■ ■ ■ ■ ■ ■ ■

HotWired (http://www.hotwired.com/) isn't a magazine in the same sense as its corporate sibling *Wired*, the unexpectedly successful digital-culture monthly. The paradigm breaks down once you're prompted for a password or you're following links in a cyberspace chain. Even HotWired's producers have yet to settle on the right metaphor for their own phenomenon. The current candidate is television: HotWired's content areas are called *channels*, their content *programming*, and their editors *producers*. The last term, at least, is hard to argue with; it captures much better than *editor* the job of assembling and constantly updating multimedia content.

If you had tuned into HotWired at the end of '95, when the network — um, magazine — was 14 months old, you'd have found seven channels to choose from, plus one special program. Some — such as Piazza (the open forum) and Renaissance (the arts channel) — had been around since HotWired's birth; others — like Adrenaline (sports) and Cocktail (drinks) — were still very new.

Each channel, listed together on HotWired's table of contents (see Fig. 1), has its own creative team responsible for content and design. Teams varying in size from Adrenaline's four to Cocktail's 12 function independently, much like the separate creative teams behind each TV network's programs. At this writing, staff numbered over 100, with more being hired every week. Indeed, at this point only two channels — Renaissance and World Beat (travel) — have full-fledged producers running the show. But HotWired, as it moves away from its original model of a single online magazine toward the new model of network programming, will soon hire a producer for each channel.

"The philosophy here now," says David Pritchard, the general manager of Renaissance, "is to push responsibility down to the channels as much as possible. The feeling is that we are best placed to understand the subject we're writing about, who out there is going to be interested, and what is going to be interesting to them." Pritchard, along with World Beat's Ross Van Woert, is in the vanguard of this push, being one of the first GMs of an individual channel.

Pritchard works with the producer of Renaissance, Michael Small, to formulate the channel's strategy and to propose budgets to the business side (led by president/CEO Andrew Anker and vice president/COO Beth Vanderslice). At least for now, channels needn't worry about revenue; advertising salespeople inhabit a separate department and focus on selling the HotWired brand name, far better known than the name of any HotWired channel. Ads in HotWired sell for $15,000 per month, with 16 to 18 ads running each month.

But while HotWired is paying the way for the channels, it's the channels that formulate budgets. "The dialog starts from us," says Pritchard. "It's not a question of HotWired coming to the channels and saying, 'This is what's happening, can you live with it?' It's a question of us saying, 'This is what we'd

Figure 1. The HotWired homepage lists each of the channels, or programs, accessible through the site. Each channel has its own creative team responsible for content and design.

like to do, this is what we feel is the most appropriate thing for us to do.'

"Our role," he continues, referring to the Renaissance team, "is very much about building audience. That's the focus: to produce a strong, compelling publication and reach an audience that will keep coming back." This means deciding where to spend money, and in the case of Renaissance money is now spent principally on improving the look and feel of the channel, making the interface easier to navigate, and paying for the best possible writers.

Additional funds are also directed at marketing, which Pritchard defines as "letting people who aren't yet aware of Renaissance know that it's there and may be of interest to them." While he declined to speak in detail, he did note that in addition to electronic marketing — which could take the form of ads on other Web sites — they're considering a variety of other media, including print.

HotWired runs the site on SGI Challenge S and Indy machines, plus two Pentium boxes running BSDI Unix — all with 64MB RAM. They use Apache server software, and the Web servers have 4GB Seagate drives. Storage is handled by the internal file system and one RAID array for the edit tree. They've got two T1 links, soon to be four, preferring multiple T1s over one T3 to avoid a single point of failure.

Just Like TV

Editorially, HotWired is overseen by managing editor Chip Bayers and executive editor Gary Wolf, each of whom came to HotWired from a slightly more established field (computer publishing and alternative journalism, respectively). On the day-to-day level they're responsible for HotWired as a package, but there's also lots of "back and forth with staff," as Bayers puts it. He and Wolf meet periodically with producers to review content and consider plans, but when it comes to the mechanics of running the channel, "each producer serves much as the executive producer of a television show or the executive editor of a magazine," Bayers claims. "Gary and I would then be the executives of the network or perhaps the cable system."

A producer such as Michael Small of Renaissance certainly has his hands full. Small is responsible for turning out at least one feature per day, generally two weeks before publication. (In the case of HotWired, "publication" means linking the piece to its contents page.) Many but not all Renaissance pieces are written by the inhouse creative team who bounce ideas off each other in formal and informal meetings.

"The people here are so talented and have such a broad range of interests that it's very easy for me to rely on them," Small explains. Producers in a new medium like the Web tend to be precocious and quick, bringing a range of expertise to the job. HotWired, being full of such people, is an ideal platform for the sort of collaborative creation Small saw so little of in his many years at *People* magazine, his prior employer.

"It can be hard to find the right balance between collaboration and decisiveness," admits Small. "You don't want to slow down progress by insisting on consulting everyone about everything. At the same time, whenever we have meetings where everybody can talk, and the designer can talk about content as much as the editors, really smart things come out of it. All I have to do at those meetings is make sure they don't go on too long."

Once the ideas have coalesced and Small has edited a piece, it's passed on to HotWired's assistant managing editor, Alexandra Huneeus. Huneeus works on headlines and jumplines, aiming for catchiness and consistency, qualities not always evident at HotWired until recently.

Next a piece will travel to Renaissance production assistant Mary Spicer, who first checks through the HTML markup (code for rendering the page). Once the text is clean, she begins assembling images, sound, and video clips as appropriate, often scouring a lengthy QuickTime movie for one downloadable clip. (Ten seconds of QuickTime can run to hundreds of kilobytes.)

Spicer then hands her work off to production manager Jill Atkinson, who works with "three," the designer of Renaissance, on the final appearance of each page. The process involves some traditional production procedures, such as specifying type, measuring column widths, and sizing graphics. It also includes procedures unique to Web

Figure 2. This page, the first page of the feature "Pastel Sentences," exemplifies the current phase of HotWired design, which stresses clean lines and graphic simplicity.

publishing, such as indexing color palettes.

Let's take a look at one recent result, the first page of the feature "Pastel Sentences," a collaboration between poet Alan Ginsberg and painter Francesco Clemente (see Fig. 2).

This page exemplifies the current phase of HotWired design, which stresses clean lines and graphic simplicity. Most titles are now set in a Swiss Gothic type, blended in with HotWired's signature version of Ad Lib. The HotWired logo that marks every page on the site is here, nicely reversed, and the program title ("Renaissance Workshop") simply presented.

Simplicity is the chief virtue of this screen, but it's helped by economy. The page loads surprisingly quickly, especially given the size of the graphics. Speed results from two very smart design techniques: breaking down a composition into smaller pieces and designing with a minimal color palette.

Though the "Pastel Sentences" page looks unified and continuous, it's actually composed of four separate images arranged vertically. The ad banner was supplied by the server, but the other three images were designed specifically for this page. They appear as a composition because they

share background color (100% red, 45% green, 26% blue), which is also the color of the window.

Three small images load much more quickly than one large one, especially if you're a Netscape user. (One of Netscape's innovations was speeding browser performance

Figure 3. Smaller images and restricted color palettes help pages load faster.

by increasing simultaneous connections.) And the images are smaller still because they were designed within a restricted color palette.

Take the title page. The largest image on the page (see Fig. 3) turns out to have the smallest color palette — a mere 16 colors. Even taken together, with lots of help from the ad banner, the total number of colors used on the page is 86, only a third the number in a maximal GIF file.

Thus HotWired's designers, enjoined to keep browser strain to a minimum, ease the flow by aiming for a constant, and low, product of image size and bit depth. Bigger image: fewer colors. Likewise, backgrounds are set to match the predominant color in the images' collective CLUT (color look-up table).

It may seem a small thing to balance elegance with efficiency, but that's the essence of good design for the Web. Some very beautiful sites are scarcely visited because waiting for pages to download gets old. And while functional pages are nice, their very efficient lack of style is boring and thus unlikely to cause much of a stir.

On the Fly

If HotWired offered just enticing content and well-designed Web pages, it would be a respectable feat unto itself. (It's also a feat replicated by several more recent Web publications — see "Select Web Magazines" on page 32.) But what they do is a lot more interesting than that.

Practically every site on the Web is a collection of static documents: HTML files, image files, sound files, and so on, that sit on a server waiting to be browsed. Most of the components of HotWired — such as article bodies and image files — are also static, but not a single finished page is. "All our pages are created on the fly," boasts Jeff Veen, HotWired's interface architect.

For example, the advertising banner at the head of the Ginsberg/Clemente piece was not placed there by hand

but was added by a script (local executable) whenever the page was called by a browser. The executable itself is hidden from view, but its trace is discernible in the source your browser receives:

```
<!-- start local exec -->
<a href="/Coin/Spnsrs/net_mci/index.cgi">
<img src="/Coin/Spnsrs/net_mci/banner.gif"
   width=476 height=56 alt="-  To find out about
   networkMCI Flat Rate Calling, click here.  -"></a>
<!-- end local exec -->
```

What you get from the "local exec" is a hot-linked image file, with HTML written on the fly.

"Ad banners all shift around," Veen explains, "and basically we have complete control over every document on the site. Just by tweaking variables here and there, we can rearrange stuff like that. The same goes for footers. If we were to add a new feature or take away a feature, for example, we can flip a switch and all the footers change. This makes production a lot easier."

It also gives HotWired a distinctive advantage over most sites, which hardcode ads onto Web pages. Ads may be moved, revised, or changed simply by updating a database and without rewriting or reloading any page. Thus HotWired's sales team has a lot of flexibility, which so far they've exercised with restraint.

But the possibilities are certainly interesting. For example, you could circulate ad banners in a channel based on a given set of distribution parameters. In other words, you could set the variables so that a certain ad would appear on a certain page with a certain degree of probability. Thus, you could guarantee exposure while minimizing predictability.

More interesting from the user's perspective are the on-the-fly modifications to the main page for Adrenaline, the sports channel. Every time the page is called, the server embeds a randomly selected splash screen from a directory, so the user gets a different screen every time. For now, it's just a spin of the wheel, but the process could easily be controlled. "We could base which page is generated on what you've already seen or haven't seen," Veen explains, "or on what's popular or not popular with all our users. Likewise, if we want to increase hits for a certain section, we could just tweak the randomness in its favor. It would be possible for HotWired to go in that direction with all its content."

Members Only

The most impressive case of on-the-fly page generation so far is the "What's New" page HotWired offers its roughly 300,000 members. Membership is open to all for free, but it does require filling out a form, choosing a user ID, and remembering a password — tasks unwelcome to a large number of Web cruisers. (We'll get back to this point momentarily.)

"What's New" is an overview of HotWired content tailored to each individual member (see Fig. 4). It shows you a table of all new and unread content in each channel, based on a set of preferences you select from a form. This bit of wizardry involves matching data in two separate Sybase databases, one containing a record of every page each member has visited, the other storing user-defined preferences.

This information is gathered at both the front end and
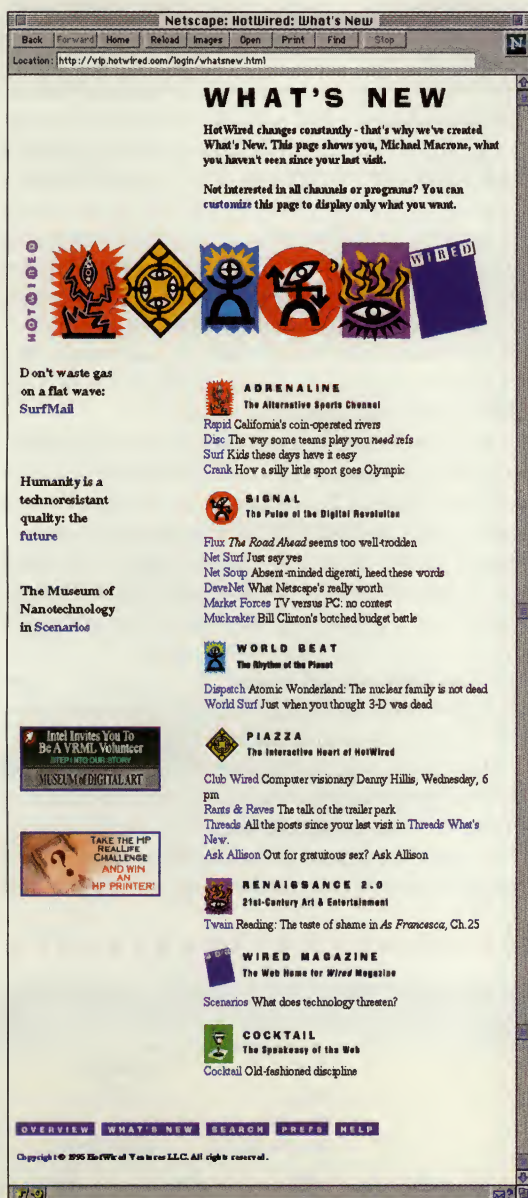
Figure 4. "What's New" is a tailored page that shows the new and unread content in each channel, based on a set of preferences you select from a form.

the back end. On the front end is the basic mechanism of HTTP authentication, supported by every server and every Web client. When you log into HotWired, your browser prompts you for a username and password; from that point until you end the session, your browser remembers that this is your authentication data for HotWired.

"Every time someone loads a URL, they have to tell us who they are," explains Veen. On the back end, a record of every hit is then added to the database. Thus, HotWired has "a profile of every user and what they've seen and haven't seen."

This data is matched with preference data to generate each user's "What's New." A variety of tools are used in the process: Perl scripts, C code, Sybase programming, whatever suits the purpose best. Similar tools are used to compose pages of user messages, known collectively as "Threads" (see Fig. 5). Users

Figure 5. Users can register preferences for how they'd like pages of user messages, known collectively as Threads, displayed.

register preferences for how they'd like Threads displayed, and that's what's generated for them on the fly.

Authenticated & Unauthenticated

Obviously, features such as "What's New" and interactive areas like Threads are possible only with those two databases — in other words, possible only for members. From the very beginning, membership in HotWired has been a controversial issue. It isn't just that registering and authentication are a pain. There were also lots of people on the Net who just didn't like the whole idea — it smacked too perceptibly of a marketing ploy.

Registration, along with HotWired's frankly commercial intentions, alienated a certain segment of the Net, though not a terribly large one. Things have only gotten worse as the Web has grown more mainstream. Not that new
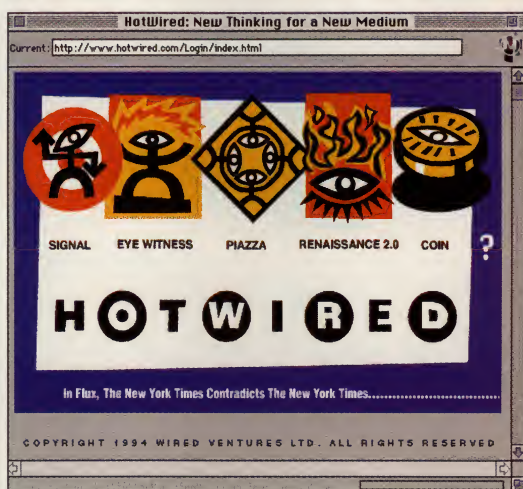
Figure 6. HotWired's original front door to the Web was graphically quite elegant, but it suffered from a certain kind of willful obscurity that many found puzzling.

From the very beginning, membership in HotWired has been a controversial issue. It isn't just that registering and authentication are a pain. There were also lots of people on the the Net who just didn't like the whole idea.

users are alienated; more often, they don't grasp the concept. For these and other reasons (such as making HotWired's pages more linkable), mandatory registration was dropped in August of '95. (You still have to join, though, if you're interested in customized pages or interactive content like chat.)

Creation of separate authenticated and unauthenticated spaces is just one way in which HotWired has adapted to experience and learned from mistakes. The most visible adaptations have been in HotWired's main page, its front

Figure 7a. HotWired's main screen, February 1995.

Figure 7b. Main screen, May 1995.

door to the Web. At the beginning, this page was graphically quite elegant, easily among the best-designed pages on the Web (see Fig. 6). But it suffered from a certain kind of willful obscurity that many found puzzling. For example, most links were labeled with highly associative names such as "Signal," or hidden under stylized, rather incommunicative icons.

Like *Wired* before it, HotWired had made certain, not 100% accurate, assumptions about its audience's tolerance for a cutting-edge interface. "One of the things we were wrong about," Jeff Veen admits, "was in thinking people would come and want to enter this world called HotWired and just explore it all. In fact, I think it made a lot of people uneasy to come in and really have no clue how much is behind all these icons and what they all mean. We assumed there would be enough information implied in the icon and in a name like Signal."

User uneasiness prompted a series of homepage revisions, first in February '95, when HotWired was four months old (see Fig. 7). You'll notice the interface is much more explicit: Everything is clearly spelled out and explained. But this too had its problems, being "so functional that it's overwhelming," to quote Veen. So in May it was back to a more elegant and understated style, which gave way finally to yet another more easily parsed variation. "The pendulum doesn't swing as far every time," notes Veen, "but we just keep going back and forth between form and function, looking for the best mix."

Balancing clarity and elegance, the all-important cool factor is one of the special challenges of the new medium. Of course, it's more difficult when what used to be "intuitive" grows strange and unfamiliar. Once upon a time it was not absurd to suppose that users would know how

Figure 7c. Current main screen.

to open a telnet window, which you still have to do to reach Club Wired, HotWired's chat area.

Jeff Veen: "People are growing less and less experienced with things like telnet as online service providers get bigger and bigger. When America Online puts four million people on the Web, you start realizing that you're no longer dealing with a Net user but maybe just a general consumer — somebody who would have no idea how to get into Club Wired. I mean no basic idea at all. We get email from users that can't cut and paste a URL. The sorts of assumptions you once could make are going away."

Before we get too caught up in nostalgia for the good old days, it's worth noting that there has been some progress. When HotWired debuted, for example, 28.8kbps modems were still new to the market, and 14.4 was barely standard. (There were even some people still surfing at 2400 baud.) Especially in hindsight, it's hard to imagine how anyone had the patience, but with the release of Netscape, things became more tolerable.

HotWired's success would have been unthinkable without Netscape, both because of the client's more efficient downloading scheme and because of its added hooks for page design. But HotWired has never required Netscape and has set standards of its own for keeping pages manageable, meaning quick to appear on a viewer's screen.

In the beginning, for example, HotWired's designers, led by Barbara Kuhr, established a palette of eight colors (eight, not eight-bit) for use in all graphics, including icons, backgrounds, and illustrations. HotWired has loosened this standard a little, relaxing the palette to four bits, but each channel's designer is still directed to make pages as simple to download and digest as possible. HotWired continues to learn more about how simplicity works on the Web and explore how to implement it. The reason, according to Veen, is "simply because I don't think people have the patience on the Web. I think you have to earn every click, as opposed to television, which people will just leave on for the sake of being on. There are so many options on the Web as well as so many technical limitations. You know, 'I've got a 14.4 modem and I'm gonna surf for an hour; stuff's gotta happen, I gotta get a payoff right away.'"

Staying Power & Java

People who spend any significant time on the Web tend to grow jaded pretty quickly, so to keep their attention you need to keep innovating. This means more than refreshing the content, adding channels, or tweaking the interface. It also means playing with new technology.

HotWired has consistently explored every significant new development in Web authoring, though they have rejected some (e.g., Netscape's server push capabilities) as too bandwidth intensive and others (e.g., Virtual Reality Modeling Language, or VRML) as too undeveloped. Right now the interface and software teams (HotWired builds most of its own software) are looking into such potential enhancements as Adobe's Acrobat (PDF) file format (http://www.adobe.com/Acrobat/Acrobat0. html) and Macromedia's Shockwave, a Web/Director interface (http://www.macromedia.com/Tools/Shockwave/).

Whether or not HotWired incorporates something like PDF depends on many things: whether it becomes broadly accepted by authors, whether there's enough bandwidth to support it, whether they can do something cool with

Figure 8. The main Cocktail screen (above) and the same screen (below) after clicking on "The Alchemist" in the left frame.

it, and so forth. The one tool that's well on its way to passing these tests is Java, a programming language (akin to C++) developed specifically for the Web by Sun Microsystems. Java applets — small applications built for specific tasks — are compact and platform independent, or rather dependent only on whether your platform's browsers support the language. (Netscape 2.0 does.) Web developers will be able to link applets to HTML documents, effectively turning the browser into an ad-hoc application window. (See "Building an Internet Web Site," *InterActivity*, March '96 for more details on programming in Java.)

One task HotWired has in mind for Java is to simplify the Club Wired interface. Where now you've got to launch an external telnet application to reach the Club, Java will bring the Club into your browser. Clickable buttons and menus will replace the obscure ASCII commands of the telnet interface, making selecting a chat channel much more intuitive.

Primarily, though, and despite exploring new tools like Java, what HotWired does to keep the interface dynamic is, as Veen says, "just pound and pound on HTML and CGI to make them do things they were never intended to do." (CGI, the Common Gateway Interface, is a protocol for running scripts over the Web.) In this, HotWired has

"I don't think people have the patience on the Web. I think you have to earn every click, as opposed to television, which people will just leave on for the sake of being on."

— Jeff Veen

had help from Netscape, which keeps extending HTML on its own initiative, in ways many Webheads will tell you God never intended.

One such new extension is frames, introduced with the first beta of Netscape Navigator 2.0. HotWired uses frames nicely on the new Cocktail channel, allowing users to click an imagemap in the left frame and see the results on the right (see Fig. 8).

The code for frames is more complicated even than that for tables, a proposed HTML extension that's become common with a push from Netscape. But Cocktail's markup is relatively simple and thus a good place to start. Here's an excerpt from the source code:

```
<HTML>
<HEAD>
<TITLE>HotWired: Cocktail</TITLE>
</HEAD>

<frameset rows="56,*">
<frame src="/cocktail/95/47/cock.html"
   marginwidth=0 marginheight=0 noresize
   scrolling="no" name="header">

<frameset cols="205,*">
<frame marginwidth=0 marginheight=2
   src="/cocktail/95/47/drink.main.html"
   noresize scrolling="no" name="main">
<frame marginwidth=0 src="/cocktail/
   95/47/drink.o.week.html"
   name="sidebar">

</frameset>
</frameset>

<noframes>

<body bgcolor=#65A902>

<!-- start local exec -->
<!-- no ad banner on this page -->
<!-- end local exec -->

<a href="/help/cocktail/index.html">
<img border=0
   src="/cocktail/95/47/stuff/cock.gif"
   height=50 width=480 alt="Cocktail: on
   the menu"></a>
```

Anyone familiar with coding tables will note some similarities in the frames syntax. For example, frames are enclosed in a larger structure called a *frameset*, just as cells are enclosed in a table. Likewise, you can control measurements for frames just as for table cells. This is the function of the tag

```
<frameset rows="56,*">
```

which specifies a depth of 56 pixels for the first row of the frameset and an arbitrary depth for the second row. (An asterisk tells the browser to resize the frame to fit the user's window.)

But framesets differ in many ways from tables, most obviously in taking URLs, rather than data, as contents.

In other words, a frameset is both an element of one HTML document (the source or parent) and a container for others; it's not an array of data but rather an array of Web documents. (You could even write a frameset to present on one screen your six favorite Web sites. Not that it's a smart idea, given your average monitor.)

When you include data in a table, the code looks something like this:

```
<TABLE>    <!-- opens the table -->

  <TR> <!-- opens the first table row -->

    <TD WIDTH=200>Some text</TD>
      <!-- the first cell, containing plain
        text -->

    <TD><IMG SRC="image.gif"></TD>
      <!-- the second cell, containing
        an image -->

  </TR> <!-- closes the row -->

</TABLE>    <!-- closes the table -->
```

But in assembling a frameset, you write code like this:

```
<FRAMESET ROWS="100,*" COLS="200,*">

  <FRAME SRC="file1.html">
    <!-- the upper left frame -->

  <FRAME SRC="file2.html">
    <!-- the upper right frame -->

  <FRAME SRC="file3.html">
    <!-- the lower left frame -->

  <FRAME SRC="file4.html">
    <!-- the lower right frame -->

</FRAMESET>
```

This code arranges four distinct URLs in an array of four frames, with a first row depth of 100 pixels and a first column width of 200 pixels. (Measurements may also be given in percentages of window size.)

Cocktail's code is slightly more involved. For one thing, it embeds a frameset within a frameset. The outer frameset contains two rows, the first occupied by a single frame and the second occupied by a two-column frameset. In schematic form:

```
<FRAMESET ROWS="56,*">       <!-- opens
   two-row frameset -->

  <FRAME><!-- content of the first row -->

  <FRAMESET COLS="205,*">
     <!-- embedded frameset: content of
        the second row -->

    <FRAME>  <!-- first column of the
      embedded frameset -->
```

```
<FRAME> <!-- second column of the
    embedded frameset -->

</FRAMESET> <!-- closes embedded
    frameset -->

</FRAMESET> <!-- closes outer frameset -->
```

This produces a browser window whose structure looks like that shown in Fig. 9.

HotWired does a few other fancy things with the basic code, such as defining frame margins and specifying whether each frame allows resizing or scroll bars. For more
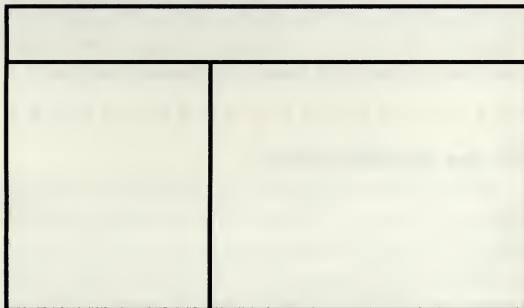
Figure 9. Cocktail uses Netscape 2.0's frames facility to produce this window structure.

on the various possible frame attributes, consult Netscape's documentation at http://home.netscape.com/assist/net_sites/frames.html.

More interesting is how the frames function when you click in them. Since framesets are arrays of documents — in effect, multiple windows collapsed into one — it's possible to refresh one frame without disturbing the others. This allows authors to have a link chosen in one frame and load a document into another (see Fig. 10). Alternatively, links can overwrite the existing (or "parent") window or spawn a new one.

What happens is determined by how frames are named and links defined. Netscape has defined a new anchor attribute, TARGET, to specify where URLs will be loaded. It looks something like this:
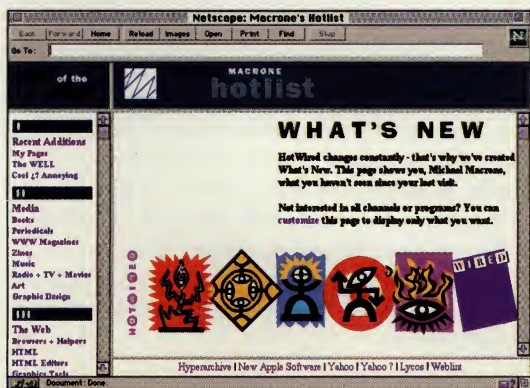
Figure 10. Framesets let authors choose a link in one frame and load a separate document into another.

```
<A HREF="link.html" TARGET="frame2">
    link</A>
```

In this case, when you click the word "link" the document "link.html" will be loaded in the frame named "frame2." Frame names are defined this way:

```
<FRAMESET COLS="40%,60%">

    <FRAME SRC="source1.html" NAME="frame1">

    <FRAME SRC="source2.html" NAME="frame2">

</FRAMESET>
```

So if "source1.html," displayed in the first frame, contains the anchor defined as

```
<A HREF="link.html" TARGET="frame2">
    link</A>
```

then clicking "link" will load "source2.html" into the adjacent frame ("frame2").

For a real life example, take the content of Cocktail's left frame, which renders this source (excerpted):

```
<HTML>

<HEAD>
<TITLE>HotWired: Cocktail</TITLE>
</HEAD>

<body bgcolor=#65A902>
<base target="sidebar">

<a href="/cgi-bin/users/imagemap/cocktail/
    95/47/stuff/old.fashioned.map">
<img src="/cocktail/95/47/stuff/old.fashioned.
    gif" ISMAP border=0
    alt="Cocktail navigation map"></a>

</BODY>
</HTML>
```

Notice the *base* tag on the sixth line: It defines the default target window (in this case, the *sidebar* frame) for displaying linked documents. Since there is only one anchor in the document, the same effect could be achieved by rewriting the seventh line:

```
<A TARGET="sidebar" HREF="/cgi-bin/
    users/imagemap/cocktail/95/47/stuff/
    old.fashioned.map">
```

In either case, the link is a server-side imagemap ("mojito.map," built on the image file "mojito.gif"). Clicking on the map sends a lookup request to the server, which sends back the appropriate document for rendering in the *sidebar* frame.

Frames would be unthinkable, except as a toy, had Netscape not built in backward compatibility for benighted users of browsers other than Navigator 2.0. This is provided by the *NOFRAMES* tag, which encloses standard HTML in case the frameset tags aren't parsable. For example:

```
<HTML>

<FRAMESET ROWS="200,300">

    <FRAME SRC="source1.html" NAME="row1">

    <FRAME SRC="source2.html" NAME="row2">

<NOFRAMES>

<BODY>

In case your browser does not support frames,
you may view my sources <A HREF= "source1.
html">here</A> and <A HREF= "source2.
html">there.</A>

</BODY>

</NOFRAMES>

</FRAMESET>
```

Everything inside the <NOFRAMES> container is ignored by a frames-capable browser; everything outside it is ignored by frames-impaired browsers. If you look back at the Cocktail frameset markup, you'll see that HotWired has used this technique so that older browsers such as NSCA Mosaic or Netscape Navigator 1.1 will still load a nice looking page (see Fig. 11).

SELECT WEB MAGAZINES

WORD
http://www.word.com/

Salon
http://www.salon1999.com/

FEED
http://www.feedmag.com/

Mediamatic
http://mmwww.xs4all.nl/Magazine/Magazine.html

Buzznet
http://www.hooked.net/buzznet/

Urban Desires
http://desires.com/

Crash Site
http://www.crashsite.com/Crash/

Suck
http://www.suck.com/

HotWired
http://www.hotwired.com/

Figure 11. The main Cocktail page as viewed with an older browser, such as NSCA Mosaic or Netscape Navigator 1.1.

On the Bleeding Edge

HotWired isn't always first to exploit new developments like Java and frames, but it's usually among the first to do it well. This has been true since the beginning, when HotWired could claim to be the first commercial online magazine, beating Time Warner's Pathfinder (http://www.pathfinder.com/) by a matter of days. (In any case, Pathfinder merely repurposed print magazine content. HotWired published original writing.)

By getting out in front of trends (advertising banners, Web-based conferencing, authenticated membership, encrypted commercial transactions, HTML extensions, Java, etc.), HotWired often influences what flourishes or dies in the ever-evolving Web medium. Along the way, it's naturally made mistakes — as in the interface design — but as usual, experience has been a very good teacher.

"Mistakes are important to make when you're developing media of any form," observes managing editor Chip Bayers. "And obviously, we learn from our mistakes ahead of everybody else. But the Web being the Web, everyone has been able to watch us through that learning process and learn from us; we've in a way functioned as a development property for every other media organization out there. They've been able to use HotWired for their R&D."

If HotWired wants to remain the market leader, it will have to live with this role — the Web being the Web. Bayers thinks HotWired still has a considerable edge, at least by computer industry standards. "We have to keep focused on the lead that we built for ourselves beginning last October, which is probably a six- to nine-month lead in any of the areas you look at: understanding the technology, building an editorial staff, figuring out how to sell advertising — all the things that go into producing a commercial Web site."

Right now HotWired seems in little danger of resting on its laurels, which would spell death in the demanding new Web market. "We just have to keep focused on maintaining our lead," Bayers repeats, incanting as much as informing, "or get sucked into the morass of the run-of-the-mill Web site."   👁

Michael Macrone designs Web pages and maintains a site at http://www.well.com/user/macrone/. He has also written numerous books on literature and the history of ideas.

NAB MultiMedia World

Co-Sponsored by:
IMA NAB BROADCASTERS

APRIL 13-18, 1996

Exhibits: APRIL 15-18, 1996

SANDS EXPO & CONVENTION CENTER

LAS VEGAS, NEVADA, USA

Improve Your Interactive Edge at NAB MultiMedia World!

Stay on top of the electronic and interactive media technologies and trends that will forever change the way we communicate and do business.

Learn how to dramatically increase your business potential through multimedia at the NAB MultiMedia World Conference. Developed and presented by industry experts and seasoned interactive media practitioners, this is one of the industry's most comprehensive programs.

Plus, with an expanded exhibition showcasing more than 150 of the world's leading multimedia manufacturers and developers you can experience hands-on the products that are revolutionizing electronic media!

REVOLUTIONIZING THE ELECTRONIC MEDIA

Digital Video

Training/Distance Learning

Enterprise Solutions

Interactive TV

Animation & Graphics

On-line Services & World Wide Web

Why Can't Johnny Ship?

BY WILLIAM VOLK

We've all heard stories about interactive titles that shipped late. How late? A few months, a year, two years, four years. . . . ■ Four years! Yes, it really happens. In fact, it happens all the time. The *Wall Street Journal* (June 12, 1995) named three major CD-ROM titles that were roughly one year late, and one title, Lucas Arts' *The Dig*, that had been

Why Can't Johnny Ship?

in production for more than four years. I know of worse examples, including outright cancellations after considerable investment — but to be kind, I won't mention any names.

These late projects often start with the best intentions. Talented people were hired or assigned to each of the tasks. The title concept was exciting and showed promise. Managers with solid track records were assigned the role of project lead. In some cases, technical risk was minimized by using a commercial authoring tool. No matter. They were still late, disappointing, or both.

Our industry has become complacent with this situation. There's a feeling that it can't improve. You can read how impossible it is to deliver products in a timely fashion in any multimedia magazine you choose. To summarize the arguments I've heard:

1. It simply takes 18 months or longer to take a title from concept to master.

Why? Is there something special about 18 months? Is it somehow built into the fabric of space/time?

2. Rewriting a title during production is necessary in order to get it right.

Did Boeing rebuild the 777 a few times to get it right? Do we tear down buildings until we get them right? Are our projects more complicated than building a new aircraft or a large building?

3. It's not possible to design an interactive product completely.

Is there something about interactive multimedia that prevents us from designing it at the same level of detail as any other kind of product?

4. We can't predict how long it will take to program a title. We don't know how long it will take to produce the content. We don't know what technological problems we'll encounter.

Obviously. What to do, then? The typical management response to the extremely high rate of failure among production teams is akin to the credit sequence in the movie *Monty Python's Holy Grail*: "The people who sacked the people who sacked the people who did the last sequence have now been sacked." So easy to fix blame. So hard to fix the problem.

The Ghost of Interactive Titles Past

Things haven't always been this way. When I published my first commercial game in 1980, it required the talents of one person: me. Back then, the design, art, programming, testing, and documentation of an interactive product was accomplished by a single person, a lone wolf. A few months or even weeks is all it took.

By 1985 I had an artist working with me and people to play-test the game. Then musicians were added to the mix. Better platforms allowed for higher production quality and more ambitious titles. Soon the staff included videographers, actors, screenwriters, designers, storyboard artists, 3D artists, animators, sound designers, producers, associate producers, programmers, system engineers, caterers, and on and on.

It's hard to remember just how ad hoc the development process was in the early days. Often you would start with a vague idea of the kind of game you wanted, write some graphics code, build some art, and play the game. Then you'd go back and rewrite the code, rework graphics, and replay the game. When staffs were small and localized, this iterative process yielded great results. It was fun to boot.

It is this history of development on the fly that haunts us with incomplete designs, inaccurate schedules, and in-the-bargain-bin products. And as a result, the creative fun of developing interactive development has been replaced by a death march: ship or die. We suffer, our products suffer, and ultimately our profits suffer.

These problems haven't gone unnoticed by the rest of the world. To quote Steven L. Eskenazi, principal at the investment bank of Alex, Brown & Sons, "The problem with this business is that companies lack a disciplined approach to product development. Management allows techie developers to start programming without a finished script. Since this is a branched medium, every time a developer makes even a slight change, it causes a rippling effect on the entire title."

Looking into the Void

There remains the possibility that we can't concoct a complete design for an interactive title. After all, the technology is constantly changing, and interactive multimedia as a genre has yet to reach any stable definition that would allow us to build products in a predictable fashion.

This view is expressed rather eloquently by Michael Moon of the market research firm Gistics. "Of the many challenges faced by the studio producing interactive works," he says, "three are most pertinent today.

"As individuals, we lack a clear understanding of what makes a great interactive title. The interactive producer today is like the film director in 1923. We're still inventing cameras and making fundamental discoveries about film and sound.

"Furthermore, interactive multimedia has yet to find its own unique creative and production process. Developing interactive titles combines the complexity of software development, which benefits from small, tightly focused teams, with the complexity of a feature film production, which supports iterative layers of work by numerous highly skilled specialists. The domains of software and film pose very different requirements for managing work. Software demands intensive communication among the core programmers, reflecting the fact that one misplaced bit can cause the entire program to crash. Film emphasizes product over process, requiring good project coordination, traffic management, and resource management.

"Finally, the market for interactive products is new and relatively immature. Producers in the Hollywood sense understand the delicate interplay between creative expression and bottom-line profitability. They're able to optimize the development of a given product to reflect its intended market. For the most part, multimedia publishers lack this degree of experience and a market in which they can apply it."

Moon makes some excellent points. However, I question the equation of film c. 1923 and multimedia today. Are the principles that underlie interactive presentation still so thoroughly uncodified? Is every title a completely new invention? Are we still making fundamental discoveries about tools and processes?

Let's look at a couple of interactive titles starting with Brøderbund's successful Living Books line of CD-ROMs. These interactive children's books present stories in text and animation that can be viewed in a linear or interactive fashion. Brøderbund created a Living Books engine that has served as the foundation for the entire product line. Although newer releases feature better animation and interactive design, the basic technology is much the same today as it was when *Just Grandma and Me* was published in 1992. The differences largely have been creative, and the creative emphasis has made Living Books a reliable purchase for parents.

Take another example, Cyan's breakthrough adventure game *Myst*. The company's first title, *Manhole*, presented a magical world of characters and places to explore. This black and white HyperCard title was one of the first ever with a user interface in which objects in the

scene serve as interface elements. If you want to go through a door, you click on the door. If you want to talk to Mr. Rabbit, you click on Mr. Rabbit. It sounds simple, but it was an original idea in 1988. Cyan refined this virtual-world concept to include animation with *Cosmic Osmo* and *Spelunx*. *Myst* combined the transparent interface with evocative imagery, inviting interactions, and a great story to produce one of the best selling CD-ROMs yet.

The road from *Manhole* to *Myst* does parallel the progress of technology. Black and white illustrations became color 3D renderings. Music segments became longer, audio quality crisper. Nonetheless, the real advances were in the creative arena. Cyan created a challenging adventure game without using opaque interface elements such as a menu or inventory — no mean feat.

I once got into an argument with a successful game developer at Art Teco 1. It began something like this:

Me: Would your CD-ROM adventure have been successful if you had published it five years ago for a black and white Macintosh?

Him: No.

Me: It's not art, then.

(Huge argument follows.)

Actually, I think his title would have been fun on a black and white Mac five years ago. I also think he missed the point. Sure, you wouldn't want to ship an adventure game in black and white today. But using the latest technology doesn't guarantee success any more than shooting a film in Panavision wins you an Oscar.

The success of the title in question is a result of the experience of playing it. And what is that experience? It's the interface, puzzles, dialog, story, and — art. By art I mean more than the sounds and pictures. I mean the design of the interactive experience. The artist masters tools, but artistic expression transcends the tools, be they hammer and chisel or parallel processors.

Art is Your Master, Technology Your Servant

Emphasizing art over technology is good not only for reliability, but for the bottom line as well. Brøderbund's earnings for the third quarter of 1995 were $7 million, up 96% from the previous year. Revenues increased 40% to $36.1 million. More revealing is that CD-ROM sales in the third quarter of 1995 represented 80% of total revenues, as opposed to 40% during the third quarter of 1994. Brøderbund's technologically conservative approach seems to have made sense in the CD-ROM market.

Think of the multimedia titles you know and love. Have we really seen earth shattering changes from year to year? I've been involved in interactive production since 1979. Yes, today's titles are a far cry from the Tandy TRS-80 and Apple II games back then. Yes, we use CD-ROMs, 3D graphics, and digital video. The games themselves, however, are still based on one or more of the interactive experiences of exploration, resource allocation, action, and strategy. Creative authors take advantage of increased technological power primarily to heighten the emotional impact of their creations. In nearly 20 years, most of the evolution has been in interface design.

Which brings us to the central point: technology is the palette that interactive designers use to create experiences for other people. Technology is not a goal. The goal is to create a title that delivers the interactive experience the designer wants. Hopefully it meets the expectations of the intended audience as well.

If we believe this, we can begin to design a production process that delivers titles on time and on budget. Titles that meet our expectations for an interactive experience. Titles designed to take advantage of the expertise of the team members involved. We can begin to create the interactive guild system that will enable us to build titles in a predictable, organized fashion.

Does this mean we should stop innovating and simply build newer versions of designs that have already proven successful? Definitely not. If most PCs are capable of playing back full-motion video, we should take advantage of it. If we have a good realtime 3D engine, by all means let's use it to improve the interactive experience. But let's not forget that the ultimate value of the experiences we create lies in the creative design of our titles.

Really On Time, Really On Budget

So can we get products on time and on budget? For a start, try answering these questions before you dive into your next project.

1. Do I have the technology I need? Can I design a great experience with the stable technology at hand?

If you feel you must advance the state of the art, then treat incorporating new technologies as an ongoing R&D project separate from title production. For the purposes of title design, use the most reliable version of any technology — you can take advantage of advances later in the project, or in the next project.

2. Is the design complete? Do I have a complete list of the production elements required? Do the artists have the information they need to create what the programmers will be expecting?

Complete designs are difficult to create. No one ever saved time by skipping the work involved. All puzzles, hotspots, game elements, and so on need to be spec'ed out. The design must be complete enough that if you handed the project to a competent off-site development group, you would get the product you want.

If there's any question about whether something is going to work, prototype it. Play on paper. Play with simple artwork. Make the mistakes before you finalize your plans. In our industry, there's entirely too much "and then a miracle occurs" between design and implementation. Allocate sufficient time to do a complete job, then *freeze the design.*

3. How long will it take to create a background, five seconds of animation, a spoken dialog, a 10 second video clip, a 3D model of a room? How much will it cost?

Most of what passes for scheduling and budgeting in our industry is guesswork. Base your time and cost analyses on experience with your creative production teams. Allow for holidays, sick days, and other periods of unproductivity. Expecting two backgrounds per day per artist isn't realistic if your artists are averaging one per day.

4. Do the programming, art, and marketing people think this design is a winner? Do the programmers believe that they can program it?

All parties concerned should participate in the design process. That means programmers, artists, designers, testers, even marketing. The goal is to arrive at a consensus that this design is great.

5. When will the programmers receive media assets? What happens if we need to move the title to a new platform?

Once production starts, generate placeholder artwork (typically rough sketches) as quickly as possible. Then track the work flow carefully to make sure media assets get to the programming staff in a steady stream. Large projects should have a librarian to keep track of assets. Don't neglect to create them in a high-quality format (24-bit graphics, 16-bit audio) for future use in a higher-quality environment.

6. Can the team produce video, animation, music, artwork, and writing on time?

Professionals cost more, but they get stuff done faster. Consider hiring seasoned veterans for production of assets and dialog.

7. Can we sustain the pace of development?

Why Can't Johnny Ship?

You're running a marathon, not a sprint. Base your production process and scheduling on a sustainable work week, not on all-nighters that extend into a murky future. A bit of a crunch near the end is typical, but if you find that your programmers have been working 65-hour weeks for the past few months, consider yourself in trouble. Deep trouble.

8. Will we need to change this design?

Changes should be avoided. If you've done the paper-play and prototypes, you shouldn't have too many surprises. If people want new features, let them know what the costs will be, especially in terms of time. Play-test a prototype of the title (or section) with changes before proceeding.

9. Is the production team being well served by management? That is, how easy is it to purchase a replacement hard disk at 3 p.m. on a Friday?

Yes, it's important to stick to a budget. Having a properly organized production process will help. However, there is a need to appear, and actually to be, responsive to the demands of the development team. Making it a bureaucratic nightmare to get that hard disk is an adversarial way of enforcing fiscal responsibility. It undermines team spirit. Remember, the big budget buster is the time spent in development beyond the scheduled ship date — not small cash outlays for necessary equipment.

10. What happens if the schedule begins to slip?

A late project is not a moral issue. Keep the team positive and focused. Deal with slippage by obtaining more of whatever you're lacking.

Have technical firefighters ready. Be prepared to use an outside production house — after all, your design is sufficiently complete to enable competent outside help to create the product you expect.

I've had wonderful experiences over the past year working with a great team of artists, designers, programmers, and other professionals in a well-organized creative and production process that emphasizes creative design over technology. It's not perfect, but I can assure you that it works. It seems impossible, I know. But, given the inherent complexity of creating interactive multimedia, it's the only sane way to proceed. I, for one, look forward to a future in which interactive titles shine with obvious, compelling brilliance because the people who make them are master artists who really know how to get the job done. ❦

HIGH-TECH GUIDES FOR A NEW-TECH WORLD

If "being digital" means keeping your head above water long enough to figure out where you fit into this new-tech world, we've got just the help you need. Our Interactive Handbook Series begins with three publications tailored to multimedia industry professionals, academia, students and consumers who wish to learn more about the complex process of creating interactive media applications. Each book breaks down one key production element, outlining in layman's terms the "how-to" procedures unique to the book's focus. Each publication is packed with examples, interviews, case studies, diagrams, business models & glossaries—— everything you need to grow your mind.

The Interactive Writer's Handbook, Interactive Music Handbook & Interactive Publisher's Handbook are garnering rave reviews. But don't take our word for it. Examine free chapters on our Web Site or call us at 1-800-529-3501 or simply visit your favorite bookstore and thumb through the pages. Whatever you do, don't get swept away by the information revolution. Grab an Interactive Handbook and ride into the "next world."

PANIC IN THE PARK

Most developers only dream of it. Others think they'll find it some day. But the crew of Chicago-based Imagination Pilots have discovered one of the holiest grails of interactive media development: the secret to designing and producing a title complete with Hollywood production values, tons of video, a rich explorable 3D world, puzzles and games, even plot, from start to finish in less than two years. A lot less than two years. ■ Even the most experienced development teams believe 18 months is the minimum time required to build a quality project. A few more adventurous souls will grant that it might be possible to grind that down to a year. But Imagination Pilots built their second CD-ROM, the interactive sequel to MGM/UA's film *Blown Away*, in six and a half months. ■ Fluke? Modern miracle? Gigantic budget, enormous staff, cut corners? Hardly. Their third title, the three-disc *Panic in the Park*, an original interactive drama starring Erika Eleniak (*Bay Watch, ET, Under Siege, The Beverly Hillbillies*), was designed and built in eight months by a full-time team of only 10 people. That same group who put out a series of CD-ROMs based on

PANIC IN THE PARK

the popular *Where's Waldo* books is also working on a virtual Erector Set and has even more Hollywood tie-ins in the works — they're two years old, and they've shipped three products with a fourth about to come out.

How do they do it?

•　　•　　•　　•　　•

On the eve of *Panic in the Park*'s commercial release, a publicist representing WarnerActive showed up at our doorstep with Imagination Pilots president and CEO Howard Tullman in tow. We had never heard of Imagination Pilots before.

Tullman, sporting a ponytail, designer glasses, and a dapper scarf draped around his neck in April, looked more high-powered Hollywood producer than nerdy game developer. Learning that IP had developed the interactive sequel to

Blown Away and that Erika Eleniak was the star of their latest interactive drama did little to dispel the impression. While visions of repurposed Hollywood tripe danced in my head, Tullman launched into his spiel.

The premise of *Panic* is that twin sisters played by Eleniak have inherited an amusement park, but the deed of ownership is missing. Someone is plotting the park's demise. It's your job to wander around the concourses, master the simulated games (think ski ball, gopher gong, and all the arcade games you played as a kid) and win tokens to play yet more games, encounter the hucksters who populate the midway (think Elvis impersonators and the snitch from *Starsky & Hutch*), gather clues to deter-

mine who the bad guy is, and rescue the park from certain ruin by finding the lost deed.

The games are modeled to respond to real physical action, so how you move the mouse directly affects how, for example, the darts fly in the balloon popping game. Tullman went on to explain that *Panic* includes some 138 alternate endings with thousands of variations, all determined by the player's actions. Its logic engine was born in the land of corporate computer-based training. In fact, using that engine for large corporate training projects is one way IP pays the rent. Which makes sense. Tullman and IP CFO Dan Di Caro built and sold a multimillion-dollar business specializing in information management before launching Imagination Pilots a few years ago.

Despite the Hollywood trappings, Tullman is refreshingly down to earth. He's eager to share his knowledge and experience. His confidence and energy give you the feeling he doesn't make a move that isn't directed at getting things done and that he usually succeeds at whatever it is he's doing.

It was six months after our first meeting, when we landed in the Imagination Pilots office on Chicago's near-northside, that we learned just how successful at getting things done Tullman and company were. IP occupies one small corner of a multistoried building that houses a furniture store and the Windy City's Planet Hollywood. Tullman owns the building.

The first thing you notice after getting buzzed into the Imagination Pilots space is the art on the walls. There's tons of it. Photorealism, modern florals that aren't kitschy, large canvases of all sorts, and press clippings about IP games cover the reception area. The phones ring constantly. People wander in and out, obviously on their way to getting something done.

Brenda Rowland, Tullman's longtime assistant, takes phone calls, fields questions from staffers, and greets me in what seems like one seamless conversation. Check off IP secret weapon number one. Rowland is the keeper of
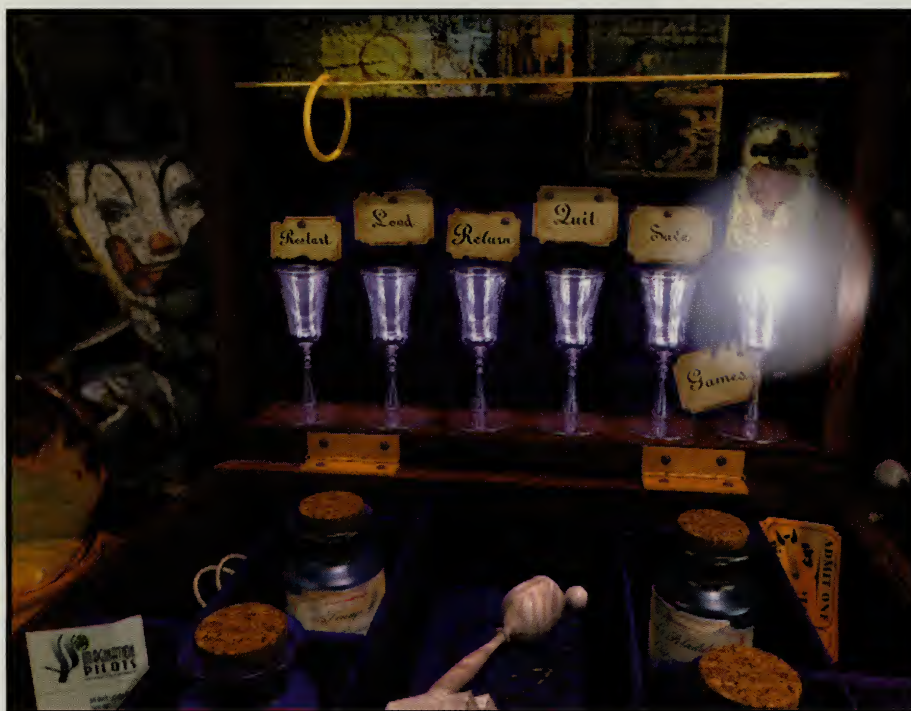
The interface for *Panic in the Park* that allows users to save, quit, restart, or load a saved game. Clicking on the corresponding goblet triggers the desired action. In the rest of the game, actions are initiated through the cursor, which changes shape to indicate when and what type of action (go left, right, play a game, etc.) is possible.

the gate, and no one is going to ruffle her organization without her say so.

On this fall '95 morning, she apologizes, explaining, "the CEO is having breakfast with the Vice President, but he'll be back before noon." As I wander down a hall to meet the rest of the IP staff, creative director Doug Heinlein refers to Tullman as the CEO, and I think, titles are a weird way to refer to people you work with. When Tullman finally pops in, interrupting my interview with Heinlein and 3D modeler Tom Miecznacowski, he apologizes for being late, explains that he'll just be a minute changing out of his business suit, and then starts talking about ideas the Vice President has about the Internet. Only then does it become clear that he's not talking about some Imagination Pilots VP. He's just had breakfast with the Vice President of the United States, Al Gore.

Tullman gets more interesting by the minute. As we settle into his office — more a loft space than an office in the traditional sense — he takes a call from his broker and I check out more worldclass art on the walls, on the floor, and anywhere else there's a blank space. Even the desk looks like it belongs in the Art Institute. But we're not there to talk art, we're there to talk interactive media design, which Tullman does as soon as he hangs up the phone.

Tullman: My view is that *Panic in the Park* goes as far as the user interface can go with point, click, and watch. The next generation has to be something that gathers more information from the user so that when we tell them the game is reacting to their input, suggestions, personality, whatever, it's not just a flagrant lie. Or when we tell the user that they have some control, it's not picking from three different logic lines and simply branching. Either behavioral engines or speech commands or various things in combination are going to be essential and that's what we're working on for future releases.

What about storytelling?

Storytelling is about suspense — there's no suspense in something you control. There is suspense in *Panic* because you don't really control it. Once you let the person drive the train, if you don't let them truly have any input, it's mislead-

ing. We see that manifest in the technology of our games. As you'll see in the demo, when you roll that ball down the alley, you can throw a gutter ball, you can do whatever you want. That's not somebody saying to you, "learn to play this game and then there's five pre-programmed routes for the ball and that's all you can do." Those are tests of complexity that over time we think the more intelligent consumer will understand. I don't know that 90% of our viewers or anybody out there looking at these things will even understand that they are materially different in kind from what is the ordinary course.

Can you describe how you designed the branching for Panic?

My objective was to create a substantial amount of variability, so what you do is an exercise in multiplication. I had a story I wanted

to tell, which was this nostalgic tale of the passage of time.

The reason I say that is because the only way the multimedia business will become real is to appeal to the principal software consumer — who is 95% male and between 31 and 45. The
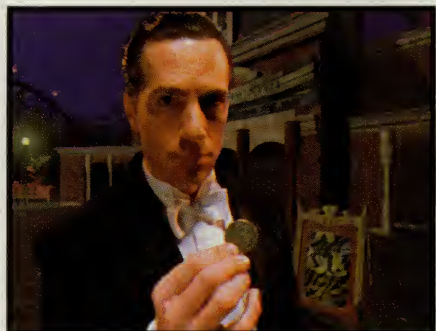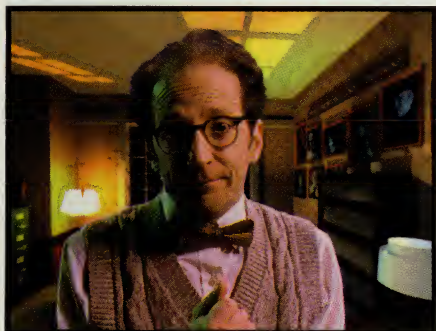
"Interactive CD-ROM today is a head and shoulders game," says Howard Tullman. "Nobody wants to see a cavalry charge where everyone is a half an inch on the screen." Shown here are a few of the *Panic* characters, including the game's star Erika Eleniak (at top, center).

![PANIC IN THE PARK](title graphic)

game kids are going to grow into multimedia consumers, but that's 10 years away.

I wanted a story that had nostalgia, characters, some complexity, but just as in *Blown Away*, I wanted it to be possible to go in and just play the games and have fun. I've always felt that *Myst*-like games require you to invest substantial amounts of time before you get anything back. They just weren't going to jump off the computer screen and come alive for the

kind of player we're targeting. I thought it was important for our titles to be easy to play but hard to master.

With this history of amusement parks, it didn't take long to decide whether the park was going to be closed or whether it would stay open. The twins concept came out of that riddle where you don't know which twin is telling the truth — one lies all the time and one doesn't. We built that into a major juncture in the game, where you watch twins fight it out. Then you have to solve that riddle. Turning that into an interactive adventure was complex. It required a lot of logical engineering to get to that point. I wanted it to evolve to a situation where at least one of the gates to the player's success was solving that riddle.

The good/bad thing was easy to formulate. Then I started to populate the park. Each of the nine characters had a good story and a bad story. The good story was related to how they came to the park and why they wanted the park to succeed. The bad story was a logical extension. They were bitter, taken advantage of, out of money, or any number of typical human conditions.

I consolidate and assemble three or four of these things simultaneously when I start the stories — motivation, character background, actual dialog — and if you look at the last 20 pages of the script, this part is just a series of comments. The questions become, "Who do I assign that comment to given that the comment is a philosophy," or whatever.

So you have these nine people, and you have good and bad scenarios for each of them. That gives you 81 variations. Then we decided there needed to be four or five other variables, so I constructed the concept of the stolen deed. The deed needed to be someplace, and it had been hidden away by the bad guy.

So now there were six locations where the deed could be. In any case, there would be a bad guy plus a location, and we added the object hidden in the Oracle's chamber and that object would be identified by the player solving a puzzle. And then you've got to go find the object in the chamber. That was a variable of which there were four or five. So now it's really simple math. You multiply nine people times six locations times five objects times

TEXTURE MAPPING an Amusement Park with Coffee Drips

Terry Schmidbauer is a freelance artist responsible for developing texture maps for *Panic in the Park*. His main tools are Photoshop 3.0 with MetaTools Kai's Power Tools and KPT Convolver, Specular Texturescape, Aldus Gallery Effects, and Fractal Painter. To create the look and feel of an amusement park, Terry and Doug Heinlein took photos of crumbling concrete, rust, and a coffee drip; scanned them; and manipulated the images as needed.

A coffee drip? "We used that same coffee drip throughout the whole game," Schmidbauer affirms, chuckling. "We did hundreds of variations on it. Coffee drips on walls — everywhere. It was great."

"Photoshop 3.0 is a RAM pig, but it's a thing of beauty," adds Heinlein. "We used its lighting effects filter a lot. We didn't have time, but we were going to build an altar to Kai Krause and have it located somewhere in the park, a devotional pyramid maybe based on a Hindu architectural motif, but we ran out of time.

"We're actually a beta site for MetaTools. We love their products," Heinlein continues. "They're excellent tools, although we had to beat on them to add the miracle inputs for vector effects — I think they didn't want to do that, but if you're in production, you have to have them."

Imagination Pilots has at least 4GB of textures in their custom library, not including 3GB of third-party collections and whatever came with the materials editors in their 3D modeling programs. Schmidbauer adds, "I don't really have a library because every project is different. You have to get what's specific to what you're working on."

IP's texture browser allows anyone on staff to access their online texture library and see at a glance what's available and track a file's size, location, creation date, file type, and so on. Seen here are 2D maps of playbills, signs, and posters.

Heinlein and Tom Miecznacowski find it easier to develop textures as they go along because as Heinlein puts it, "It always takes longer to look for the right texture map than to draw some bricks in Illustrator, decay them in Photoshop, and apply them in 3D Studio or Animation Master. Once again, if you're a good artist, you know how to do that. You're not relying on somebody else's input. We detail it down to the matchbook, which means that Terry had to figure out what a matchbook looked like and do a convincing representation of an old

four doorways and, lo and behold, you have a few hundred potential solutions.

As you encounter each character, their dialog is variable depending on whether they are the villain. If they're not, they may make comments about the villain that would give you a clue or hint.

That keeps the player honest because it means having to watch the movie each time. Sometimes the new information is at the end of a clip, and sometimes it's at the beginning. Sometimes the clip says nothing about the situation. To that equation, we added these four concourses and a guide to each concourse who is there to assist you. These are complex geographical areas. One of the issues was making sure the players didn't get lost and that they knew what was expected of them. The guides then become intelligent branching. They know whether you've played a game, whether you've

Most of the games in *Panic* are traditional arcade games. When designing the games, Tullman felt it was important that they be easy to play but hard to master.

matchbook from the '30s. I asked him to do an old playbill that is just a clue in the game. He doesn't have time to spend a day in the library looking up stage bills. He took a new stage bill and just pretended it."

Convolver didn't figure into the project until late in the game, though all admit they use it constantly since acquiring it. Says Heinlein, "I use it on almost every file."

If you're surprised that they don't keep better track of textures, yet still get their games built in record time, so were we. But Miecznacowski is adamant that it's easier for him to find a texture close to what he needs or simply make a new one.

What if you've made the texture before? "There is some crossover, but that isn't really predominant," Heinlein explains. "That may happen once in a while, but it's no worse than losing files. At one point, we lost an entire concourse when it had been rendered already, but we were either going to re-do it from scratch or live with what we had and work on the TGA files as they'd been rendered. We didn't have time to rebuild it. I mean, the geometry just got lost. We have tape backup, CD backup, and you still lose stuff. With the movie files and the sound files and the art files, you're swapping gigs in the 40s and 50s. Not only that, we have a 14 gig network, and it was out of space every third day. We were burning CDs constantly. Then you have your textures on this CD, and when you need it, it's gone because someone else has a piece of information that they need from it, and you don't know whose office the CD is in. So you just say, screw it. I don't have time to track it down so I will just build what I want."

Miecznacowski continues, "Also when you're doing a project, you keep a large texture library active so it's not a matter of having to find a texture. I had them in three places — on my shared file, on my active programs, and stuck away on a CD, too."

IP also has a file on their network that holds 3GB of textures, so anyone on staff can access their texture browser and look for something (see example on facing page). Catalogs keep track of what's where. They can track a texture; the objects it's associated with; its location, size, and creation data; and when it was last modified. Which proves very handy when the art department is called

to do a partial build for tradeshow demos well before the game is finished.

"We're also handling promo material," says Heinlein. "We have to pick shots that are exceptional views, rerendering them to 1280 pixels because we have to have material for the salespeople to take around."

A few examples of textures and signage, plus a look at how a simple coffee drip can make a metal door look wonderfully weathered.

PANIC IN THE PARK

gone in a certain area. They know a lot of things, so they give you pieces of strategic advice throughout the game. And that varies every time you encounter them, depending on what you have done.

When you do the math, there are about 5,000 different outcomes, but the principal outcomes involve where the deed is, what the object is, and who's done it. We think some players will play through that a number of times, some will want to see each of these people confess because they'll be somewhat invested in the characters and frankly, that's the second part of storytelling. One is suspense, and two — equally critical — is that you care about the characters. If you don't establish those kinds of connections, you aren't going to be successful.

I don't regard it as "content is king" as much as I regard it as "craft is essential." The way you present your characters is important. You need to figure out ways in short pieces of business or costume or makeup or dialog to establish a connection so the player is interested in hearing more from that character.

Is there anything about the player that will affect your approach?

The age of the player. Younger people aren't interested in people per se. They're interested in cartoon people. That isn't what we are trying to accomplish. Yet another population in the game — the vendors, the people who run the games in *Panic* — are all cartoons. They're abusive. They kid you. They joke. There's a whole other world that's presented. You've got the park administrators, the guides, the bad people, this *Wizard of Oz* character. You've got this little city there, which is what we were trying to create, so that's where the branching takes you.

What about the design of the arcade games?

With two or three exceptions, the games are traditional arcade games. Gopher Gong, Ski Ball, Throwing the Darts, Rock and Roll — which is taken from a kid's game where you have to tilt a board to get BBs in holes. Those were pretty standard. The most complex is Blind Man's Bluff, which I've been wanting to visualize for a while. The idea is to have a maze that you could not see. You have to work through it by trial and error faster than the computer can. The computer is equally dumb, so it also has to do trial and error, but it has a longer path than the player. So if you were equally skilled, you'd learn very quickly what was going on. You have to go left, right, up, down, whatever.

That turned out to be a level of visualization that was so tough we didn't know if the player could be successful. We took these marionettes and suspended them over the play board, so you still have to figure out a pattern. You watch a color code plus a value code and then move your marionette in the direction of that color code, which changes. When I say changes, I mean red

is north, but only sometimes. In the meantime, the computer is trying to beat you. It takes a while to develop the skill to do it. Once you develop that skill, it's very compelling because it's akin to Pacman, where you learn more and more of the route each time and then go try it again.

In *Blown Away* the games were random puzzles. In *Panic* the games are skill-based, and what's random is the total circumstances surrounding the role-playing game. But in the games, you learn to get better and better. Over time, you should be able to master all the games.

Tell us about your database technology.

The largest thing in this business is to engage in the discipline of knowing where every single asset, every single phrase, is. You might think, oh, here are a bunch of taunts this vendor says, but he's one of 14 different people. And you're going to have a tape somewhere, and you're not going to know if it's been captured, digitized, sweetened, had the music added to it or not. Multiply that by 20,000.

You have to have the discipline from day one, phase one to have every single asset identified by type, so this really is more than a script in a traditional sense. This is why very few people understand what goes into these interactive things. It's not only a script, it also contains programming notes. It contains rules, logic, everything. We also have logic charts, but writing this from the top down, it has props, it has stage directions, everything.

So the first part of the database is the concept of organizing every asset, everything in the script, to be recapturable and managed.

Then [*pulls out a thick document*] this is, by speaker or by order, all the movie assets. This is important for translation, and it also gives status descriptions and stuff like that. The other significance of the database is, again, every asset may or may not be working successfully. These are all assets that needed transitions, so every one of those was not simply an asset, it was also a physical location within an environment. So the database has to say, "Now I exist. Now I am working properly." The database feeds the assets into a comprehensive bug testing and anomaly reporting system as well.

That's all automated?

Yeah. It's all automated.

Then, this is a game logic chart, which is developed at the beginning. It literally takes you through all the alternatives, and it gives you jump information and sound information. This document eventually becomes a testing tool, but what's critical when you have a hugely variable game is to be able to test each of these things on the fly. Built into the game is a thing that says okay, I have been through here. Now I want to go through as if I were Annabel Lee, then it makes a quick change and when you go back through it, it plays differently. That's the only way you can possibly keep track of this unbelievable com-

What was originally scheduled to be seven days of bluescreen shooting on 35mm film was whittled down to five days by producer Tim Tennant and director Ken Berris to meet budget. The shoot took place at Culver City Studios in California.

plexity of testing all the alternatives. That's the logic part of the process.

Then, these are my drawings of another issue, which is, we obey literal directions. I have to determine what's fair to the player. It's not fair to have the next click suddenly turn the player 180° in the wrong direction, so we test over and over from each position — there are a lot of things inherent to that.

First there's the left or right choice. If I go this way, I meet an actor. But that actor has to know that I didn't turn around halfway, pass him, and not meet him. So he's got to be smart enough to make a relevant comment like, "Hey, you can walk all the way around the park and come back here, but. . . ."

Each of these views had to be tested to see what's the logical forward move, what happens if they go left, if they go right, so we don't make somebody do right angle crazy kinds of turns. Everything should be pretty intuitive. Then we have clues and animations buried within the situations as well.

There are six or eight different clue systems. You'll do something like pop the balloons and something will come up that suggests who did what. There's an online electronic map of the four concourses. You have a flashing light that's always available to tell you where you are, as a guide. Again, important from a player's point of view.

The database lets us grab elements. Take bugs. It lets us track a bug, who found it, when was it found, who is going to deal with it, what version (because we have version control), whether it has to be fixed, what's the nature of the thing, and what assets are involved.

From the standpoint of trying to understand how complicated it is just to turn one little clip into a thing, this [points to another pile of printouts] is a measure of a different part of the database, which says as to every single video asset, have I captured it, do I have the sound to go with

Using FoxPro as their relational database, Imagination Pilots is able to track every asset through every step of production. At any time, the IP staff can get status reports that indicate what's done, what's pending, who's responsible for producing what, and so on. The database also feeds the assets into a comprehensive bug testing and anomaly reporting system. "You have to have the discipline from day one, phase one to have every single asset identified," says IP CEO Howard Tullman. The application of their database shows up in every phase of production. For example, Tullman used it to help director Ken Berris keep track of dialog on the bluescreen set, while the art department uses it to track everything from geometry to textures, animation, and video.

it, do I need an animation, have I composited it, when is it done — spread out over thousands of assets. So that at any time we can see where we are, which is important because if we want to bring forward a preview of something, we know exactly what we have and what we need.

We track movie assets, sound assets, art assets,

and we have charts that track each arcade design's rules, scoring, programming variables, things we're testing, and then all these effects — lose, win, ball rolls, ball hits the side, what are the artwork requirements, what's linked to it.

How do you attach the filenames to all of these assets?

They all fall out of the original script, but when you add another asset, the system knows. It says, here's the last one, so it eliminates duplicates. It tests your cell.

Does the system help you name assets?

No. I created the system, and I had to keep the assets in my head, and when eventually they cranked it up, they would come back to me and say, you were 97% correct, and there were three assets that were goofs or something, and how do you want to rename them? Also, when we started shooting video, sometimes the actors would have additions, good stuff. And frankly, at one point, we didn't have good instructions for the stair game, so on the fly we had to write some dialog for the Oracle. I had to shorten some speeches because they turned out to be way long and boring. Then we had this mugger who could not say his lines. I mean, he just absolutely could not say his damn lines. So we kept shortening them. That was a real piece of work.

Then there are all the programming consid-

The online map of the park showing the four concourses, which internally were referred to as Concourses A, B, C, and D. Compare this finished game map with the working diagram on page 51 of Concourse D — a.k.a. Bug House Square — used by the design team to keep track of all the assets that needed to be generated to bring the area to life.

PANIC IN THE PARK

erations of the database, which are enormous.

In working with the actors and actresses, did you encounter any particular challenges?

There are three or four different kinds. First, we used commercial instead of theatrical talent. A very significant difference. One, they don't bring a lot of attitude. Two, they're used to absolutely one day. "Who am I today?" Boom. Tomorrow, "Who am I now?" They don't require a lot. So they were extremely grateful to have a script they could get into. They were given a little more meat than they were accustomed to. They weren't treated like robots, and they were given a chance to play more emotional things, so all of that was good.

The bluescreen issue is a terrible issue, however. It's complex. Players were materially easier to deal with when they were doing twosomes, as opposed to playing to the camera. We filmed two days in the courtroom, where they played

"The bluescreen issue is a terrible issue," says Tullman. "Players were materially easier to deal with when they were doing twosomes, as opposed to playing to the camera. . . theatrical training tells you to ignore the camera. Here we were saying focus on it."

to other actors.

Interactive CD-ROM today is a head and shoulders game. By that I mean, given the size of this tube, you had better be this close and use full-screen, full-motion video because nobody wants to see a cavalry charge where everyone

is half an inch on the screen. It's pretty convincing if someone is here talking to me, especially in our new games when it's running at 30 frames per second.

Getting them to play dead-on to the camera as if they were really talking to a person is something that even Erika picked up on. Ninety-nine percent of theatrical training tells you to ignore the camera, and here we're saying focus on it and your words are going to be in the TelePrompTer in the camera. It took forever to teach them to use that as a security blanket, but not to read.

We ended up saying, "Look, you're an alcoholic old guy, and you're pissed off. It's okay to take the line and do something with it." It didn't have to be the exact line.

That was the last thing. I am not a SAG director — so we had a SAG director — but I am very invested in these lines. Some of these people didn't hit the lines. But the director and I are very good friends. We were able to work together, so I could whisper in his ear. I would never give direction unless I was invited to, which I rarely was.

Again, our database technology was very useful because while he was working with the script and had a script girl to tell him where to go next, I was able to tell him my take on the scene. I had all the lines for that speaker on one page, and I was able to show the actors and actresses what was going on.

As far as the actors and actresses went, we had very good results. When we were trying to attach a star, we were trying to think about how much
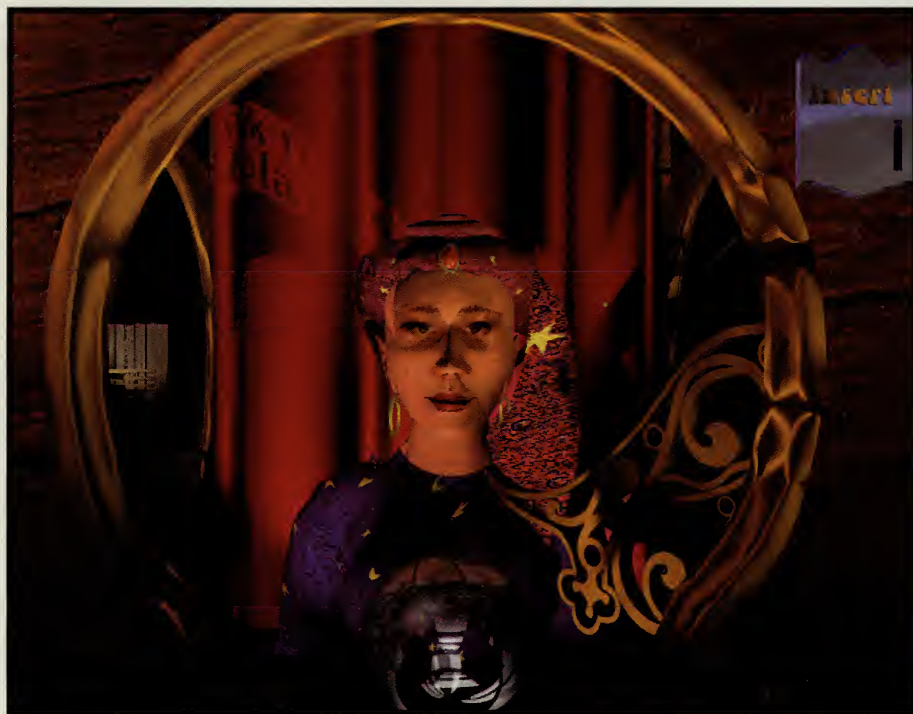
"My view is that *Panic in the Park* goes as far as the user interface can go with point, click, and watch. The next generation has to be something that gathers more information from the user so that when we tell them the game is reacting to their input, suggestions, personality, whatever, it's not just a flagrant lie." — Howard Tullman.

value a star adds, how many dollars to spend, and what kind of star to use.

We made a list: Teri Hatcher, Courteney Cox, Jennifer Aniston before *Friends*, Cameron Diaz before *The Mask*, and Erika. Teri Hatcher and Erika have astonishing cue recognition on the Internet. Erika because she's naked all over the place. Teri Hatcher less so, except for one or two bathing suit things. It's just that she is a gorgeous girl. So my thought was, that helps. So we went after a couple of people, and we ended up with Erika.

When she showed up on the set, she looked at the wardrobe. We didn't intend for our heroine to be a sleaze ball, but we thought it would be useful for her to be dressed as a carny and wear tank tops or cutoffs or something like that. I mean, this is a girl who played totally naked in several roles, but to our chagrin, when she came on the set, she decided she was a serious actress.

We had to do a completely new wardrobe. In fact, it was so crazy that we brought her a pair of jeans and a t-shirt and she said she wouldn't wear them, even though it was the stuff she had come to work in.

But overall?

We were very pleased with the cast as a whole.　　　　　　　　　　　　　　🐾

"As a manager who's hiring people," explains creative director Doug Heinlein, "I'm looking for good mechanics, but they really need to be good artists. I've got drawers full of resumés from people who may know their programs, but they aren't necessarily good artists."

Doug Heinlein & Tom Miecznacowski

ON WORKLOADS, MODELING, ASSET MANAGEMENT & DESIGN

Doug Heinlein is the creative director for Imagination Pilots. The path he followed to get there includes drawing comic books, running a storyboarding business out of his home, and teaching himself Mac-based design because there happened to be work available for people who understood Illustrator, Photoshop, and modeling programs.

A Mac sits to one side of his desk, but the dominant piece of furniture in his dimly lit office is a drafting table loaded with conceptual sketches. Outside his door is a spartan workspace shared by three modelers and a varying number of interns. One of those modelers, to my surprise, is Tom Miecznacowski, one of my best friends from high school. He's been doing 3D modeling for game developers in Chicago. (His credits include Viacom's *Are You Afraid of the Dark?* Before that he was doing broadcast animation and designing corporate slideshows.) I didn't know he was on staff at IP until that moment.

After a round of how ya been and whadda ya think about da Bulls, da Bears, and da baseball season, we settle into Heinlein's office. Surrounded by sketches, storyboards, and a mound of notebooks — laserprinted copies of the *Panic in the Park* script and production spreadsheets that represent IP's asset management system — Heinlein and Miecznacowski get right to the heart of what makes Imagination Pilots special.

Heinlein: We turn out things real fast. *Panic in the Park* was started in November '94 and it's off the shelf now. We shipped it in July '95. Not only that, we did it with a barebones staff — we have two and a half modelers and two illustrators if you count me.

One of my challenges as creative director is that I'm looking for people who

can think because I don't have time to sit down and storyboard my butt off. I'll give a set of general parameters that I need to see. Maybe it's harder for them at first, but I think Tom now knows what I'm looking for. I can just go to him and say here's the dimensions and here's the border of the concept. You fill in the middle.

As a manager who's hiring people, I'm looking for good mechanics, but they really have to be good artists. I've got drawers full of resumés from people who may be good mechanics and know their programs, but they aren't necessarily good artists.

Anyway, we did *Blown Away* in about six and a half months. We did *Panic* with a slightly larger staff, but it is a three-CD set, and we did that in about eight months. You don't have the finances now to spend two and a half years and $6 million on a product that's just not going to hold up in the marketplace. If you have to sell 30 million copies, good luck. You have to be able to get the stuff in and get it out and get it on the shelves and get on to the next title pretty darn quickly.

What's your workload like?

Panic in the Park was a quantum leap forward for us as employees in this company. It was tough on us, but it was nothing like doing *Blown Away*. We actually had some weekends off. I was here on New Year's Day, but we had parts of weekends off, you know. That's what it takes to do this kind of stuff. It's not like we're working from five in the morning to one the next morning, either. We had extended business hours regularly and at least one day a weekend for sure. But everyone has kids — I have kids, the CEO has kids, so there's an understanding that you need that part of your life.

We're not lunatics, but we do work really hard. It comes and goes. We actually

Doug Heinlein

Tom Miecznacowski

use a lot of freelance workers. At any one time I think there were like six people over there [*points to the space the modelers occupy*]. Now I have three artists.

Miecznacowski: There's never really been more than three or four. We might have an intern in here doing sprites. But at the most we have three modelers and an intern or two.

Heinlein: It's really great to use interns — it's kind of hard on us because we have to slow down to get them going, but it's great to work with them. They're really good kids.

How long do they usually stick around?

We've lost two interns who became full-time employees and then left. I took a gal straight out of college who was a modeler. She was with us just over a year. Actually, she was a full-time hire and she split. She recommended this person who was an intern for about eight weeks, came on board full-time, then decided to go into education. He wasn't really into games, so he went the educational route. We have a new 2D guy, Dave, who'd been an intern for about seven months. We just brought him on full-time. We've got a new intern but he keeps getting beat up by strange people on the streets, so I don't know if he's going to survive us or not [*laughs*]. He's been with us for at least half a day so far. He's real bright. We have an awesome intern program because they have to do stuff and we pay them. It's not just getting coffee. You get to sprite. One intern did some Director stuff for us, which was excellent.

Did he come in knowing Director already, or did you train him?

I think he knew a smattering of it from school. I basically said, "Hey Dave. You're the last man on the ship and this needs to get done. Here's the tools, I'll see you tomorrow. Have fun." It didn't come together overnight, but then no one expected it to.

Miecznacowski: Yes, you did [*laughs*].

So you said, "Here's the box, go figure out how to do Lingo code?"

Heinlein: Yeah. If you're going to be a small company and be successful in the marketplace, you have to rely on the ingenuity and creativity of your people. If they don't have the internal resources as people to meet a challenge, then they're not going to last very long here. They might last at a company that's a lot fatter and can afford to have personnel sitting around throwing basketballs into a hoop all day. I don't know what some of those other companies are like because I've never been there. We manage to have a good deal of fun. We have a whole collection of weapons here, these little ball-firing things [*picks up a kid's toy*], so we have our little wars here, but it's certainly not as fat as other places. At the same time, for me, it would be nice to grow larger. We have two projects going right now with a third in the works.

It's great because I get to learn a bunch of stuff. I came in with pretty good Photoshop skills, pretty good Fractal Painter skills, decent modeling skills. My biggest moan is that I came here to model and animate, but I have to leave that to other people. I've been managing, directing, and doing a lot of 2D work. Lots of concept stuff.

You've got a Mac Quadra on your desk, but the modelers have PCs. What software are they using?

At this point I'm actually facing the decision of getting a Power PC or buying an Intel platform. I'm leaning toward Intel, because I want to work with Lightwave. It looks like a nice program. Right now we're using 3D Studio as our primary modeling tool. Although when Tom came on board, he brought in Animation Master, which is excellent. It takes forever to render, and you can print that. Martin, speed up that rendering! But it is really a very nice program. To supplement that we just picked up Lightwave.

Miecznacowski: Lightwave is definitely going to be a hard charger. It's an

excellent program. I'm comfortable on it. Our other modelers are picking it up. It's got some real strengths compared to 3D Studio.

Such as?

Miecznacowski: Undo. The new release of 3D Studio is much improved, but we're not using that.

You're using 3D Studio release 3 instead of 4?

Heinlein: No, we're using 4. I don't know if we're going to beta test 5.0 or not. We've had discussions about it, and I hope we can be a part of that program because I would love to see it. Apparently, it has a lot of new features.

Miecznacowski: It's a ground-up rebuild, but it's the same philosophy. It's strong, but I just don't happen to care for it. I come from the broadcast side.

RENDER TO PREVIEW 3/14/95					
ASSET	PRIORITY	FRAMES	STATUS	EST> TIME	DUE
CCBM3G2	1	150	RENDERED	12	3/14/95
CCBM3G2A					
CCBM3G2B					
CCBM3G3					
CCBM3G5					
CCBM3GA					
CCBM3GB					
CCAM1A	2	50	RENDERING	6	3/14/95
CCAM1B					
CCAM1C					
CCBM1A	3	50	queued	6	3/15/95
CCBM1B					
CCBM1C					
CCAM1EVA	4	50	queued	6	3/15/95
CCAM1EVB					
CCAM1EVC					
CCBM1EV	5	100	queued	15	3/16/95
CCBM1EV					
CCBM1EV					
CCBM3EV	6	100	queued	15	3/17/95
CCFM1	7	100	queued	8	3/18/95
CXXM3		1	RENDERED		
CORM18B		1	RENDERED		
CORM4EV					
CORM4EVA					
CORM4EVB					
CORMEVC					
CMGMTV1		1	RENDERED		

By glancing at spreadsheets, such as this one showing movie assets, the IP staff can see at a glance how big a file is, its status, due date, and how long it may take to render.

I'm used to splines and NURBs, so I like Animation Master. Lightwave is a nice combination. It's got excellent modeling tools and a fast renderer.

Heinlein: A part of it is just aesthetic — it is just my choice in terms of what it feels like. I can see immediately the difference between a 3D Studio image and an image created in another program. It's a gut feeling. That brings in my painter background of warmth and atmosphere and depth of field and focus, you know all these intangible, very hard to program areas. It's an excellent program and as we move forward, it has become apparent to me that it is not a matter of what platform you're using, it's "what's the job you need to do?" and then just go do it.

If we need an explosion and we're not using Alias, I am more than happy to go to someone two weeks ahead of time and say this is what we need, you'll do that portion, and then we'll just plug it into the game as we need it. But we just did an Alias/SGI thing here, and we decided not to go with it because for us it just wasn't cost efficient. As we grow as a company, I think we will, but that's basically a seat that's so hot you need to keep that thing running 24 hours a day. Maybe in about four months, when we have 10 projects to do, then we'll do it [*laughs*].

Let's get specific and talk about how the projects develop from start to finish.

Heinlein: Panic was written by our CEO. I started with him last fall. We had this idea of what he wanted to do, and we just started talking and firing faxes back and forth. But he knew pretty early on that he wanted to do a script set in an amusement park. He came to me and said, "I have this idea for interactivity.

It will have this amusement park and in the center will be four concourses you can walk through. The four choices start at the beginning of the program. By the end, you wind up in this central area." So I said okay, fine, that's great. I haven't been to an amusement park in 20 years. So I went out and spent a day at a local amusement park, and then we had a lot of photos coming in from Coney Island and another small place in upstate New York. We started to gather our photo samples. Got books from the library. Spent a good amount of time doing research, just grabbing all the material I could grab. Basically, Howard Tullman had a memory of his youth in Chicago at Riverview [Chicago's answer to Coney Island], so he was filling in this echo from his own past. I was not born and raised here so I didn't know much about it. So there were some objects that he really wanted to bring in. I don't know why he wanted to repopulate this park, but he did, and it was my job to do it.

Miecznacowski: And it's his company. . .

Heinlein: . . . so he can do whatever he wants [*laughs*].

As the script developed, we knew we had to have our main gate and I sat down and tried to do an overview in Illustrator. What I did here [*see the map on the right*] was to start developing all the down shots, where we would be as we worked our way through this park. These maps are broken up into three sections for my management. One describes all the stills, so it's all of the camera positions. One group of maps is where all the video happens, and then one is where all the models are — trying to keep track of those. Now, looking back, if I had known it was going to be like this I would have committed suicide.

So we just started building based on my floor plans. To start off, we needed a park that was basically shaped like this. I would draw stuff in Illustrator and hand it off to a modeler and then we just worked in conjunction.

Let me just run you through a sampling of what it looks like. We have the four concourses, which for our internal uses were known as A, B, C, and D. In the script they are referred to as "The Street of Dreams," "The Miracle Mile," and so on.

We knew we were going to have video. We knew we were going to start up here and walk down this way. What I was trying to do was design a model that could be broken up and rendered in chunks, so we would not be faced with god-awful rendering times of, like, 25 hours for one frame. We were going to do hundreds and thousands of frames for all the camera positions that had to be done. Besides that, as the player navigates through the space, there's all the back-grounds for all the interactive movies that happen. So this whole area might be rendered for a total of 18 hours of screen time if you count all the backdrops and stuff.

We knew the player was going to be forced to come down this way and that here [*points to an eight-sided compass mark-ing on the map*] some decisions would be made about where to go. We knew that there were going to be three games per concourse. The whole way you get through this is, you are given the background story at the main gate as you enter the game, and then you choose one of the four concourses to start with. As you move through the concourse, you're presented with video options depending on how you navigate. You're told that you have to play these three games to get through the concourse successfully, at the end of which you are given a magic coin to be placed in an object in the gazebo area, which is the central area. When you've been through all four concourses, you're moved to the second level of the game. You have to go through Concourse A, win all the games, do the video stuff to get the back story, and do B, C, and D, at which point you go back to the gazebo, put the last coin in the machine, and solve two puzzles; a trap door opens, and you are transported down to the second level of the game where you continue your search.

Anyway, we structured the whole thing so that we could render this out, render that out, etc. One of the limitations we had with

3D Studio — it was a real problem; I've talked to the people at Autodesk and I think this problem has been cleared up with release 5 — is that they had a map limitation, so the program allowed us to run only 235 maps. We were trying to build a small city, and besides applying textures to walls, we had tons of signage to put in there. It was not a pretty thing.

Miecznacowski: There was a limit on lights too. We kept hitting the wall.

Heinlein: We were trying to figure out why we were crashing a lot in the early stages. Then it finally surfaced that there were these limitations in the program. Then it was just a matter of "oh, we'll cut back." So we started cutting back on some stuff, but it still looks great.

So that's how this was laid out. Howard went in and placed his characters on this map, so this map is in the final version, but it was just developing as we went through it. We knew we had to have these little nooks and crannies and sides of buildings for people to come out of, and doorways and places tucked away where you make a left turn in the game and this guy comes out and talks to you or something.

Miecznacowski: We usually did quick wireframe renderings that we could give to Tim [Tennant] to take on the video shoot with him, so he'd know what the line art was going to look like. That way, the scenes would match the live action real nice, so you could have the camera pan and get a nice move here and there.

Heinlein: The video was shot in California on bluescreen — at Culver City Studios I think. It was all shot out there by Ken Berris. Howard, Tim, and our

SKY PARK
PANIC in the Park LEVEL 1
Concourse D
1/11/95

One of many working maps used to keep track of production assets — in this case the models that populate Concourse D. Navigation points are represented by the eight-sided compass points. Models are indicated by their filenames, for example CCDGMOON — CCD for Concourse D; G for geometry (an M would indicate Movie; an A, animation); MOON for the model itself.

video person at that time went out with them. I think it was shot in a week. They came back with all this stuff they had to digitize. I believe they shot on 16mm. We shot on Beta for *Blown Away*. Then we went to work incorporating the characters into the backgrounds. We worked with the video department to develop backgrounds for these people.

What was it all composited in?

Heinlein: They used Avid and Premiere running a Radius Telecast on a Power Mac. When the first Power Mac 110s came out they grabbed one. I think they used a Quadra 840 and a 110 for all the compositing work. We have a rendering farm here. I don't how proprietary that is, but we have a number of Pentiums with RAM up the wazoo. The rendering farm worked almost without a hitch. We really ran into very few problems with it over the course of the project. Every now and then it would crash. The nice thing was that we're on a BBS here, so I could dial in from my home, log on, and check on the drives to see if the renderings were still coming out. So at like one in the morning I would check and make sure that the rendering was going on. If it wasn't going on then I would have to get in my car and drive for 35 minutes to come check it out or reboot it.

Miecznacowski: Most of my weekend work was coming in to restart something.

How do you guys feed the images to the rendering farm? Is it sneaker net or are you networked?

Heinlein: We're networked internally.

Miecznacowski: Yeah, 3D Studio has its own networked rendering stuff. I ran the Animation Master stuff on dedicated machines because their net rendering stuff is difficult. So I had three Pentiums running, each for a separate scene.

Heinlein: We had to keep track of where the models were and what the progress was. When it comes down to the asset management, you have assets for everything. We had still assets, sprite assets, sound and video assets, but then we had geometry assets too. Someone is given the job of doing 14 buildings and you need to know that those buildings are going to be ready when the other things are ready. The 2D people have to have the texture maps ready to go when the modeler needs them. Facilitating all that was a big chunk of my time. I ended up doing a lot of the 2D maps, although I got no credit for it [*laughs*]. We have a good asset management system that was created inhouse. We use just eight-character filenames so everything has to line up.

You can see here [*points to another printout*] that we have a CCAM, which means a Concourse A movie. Over here, CCB is Concourse B geometry, and WOR is its designation as a work trench that had to be modeled. All this stuff is geometry names so we can keep track of various pieces. From that we would assign texture maps using the same coding. That way, we knew that all of these texture maps belonged to this building, which belonged here.

Really, the biggest problem was lighting. With the inclusion of video, besides

the fact that you're building a small city that has to be texture mapped and needs to exist in a kind of realistic space, we had five pieces that needed to fit together as a whole. So early on in the project, after I had visited some of these parks, I noticed that there were these big rides. They became like geographic landmarks. I wanted to make sure we had a dome over the entire place, so we did a skydome. We used KPT Bryce to develop a sky. We had an orientation toward which the models would all face, so we knew where north and south were for everything. And then we put the dome over each model when it was rendered, so if you are facing north in Concourse A you are seeing the same sky you would see in Concourse D. Even in 3D Studio, we knew where our center was, so we could stick this concourse a little bit over here to provide a sense of depth.

I figured out that 12 objects would loom really large, like a ferris wheel, a big ride. We had a couple of big roller coasters that were really very simple models that were basically like silhouettes painted black and lit up on the horizon. I think this is one of the reasons it feels so good when you're walking around in there. You get this feeling you are in a place. It's not just thrown together haphazardly. Somebody has thought not only about the space you are currently
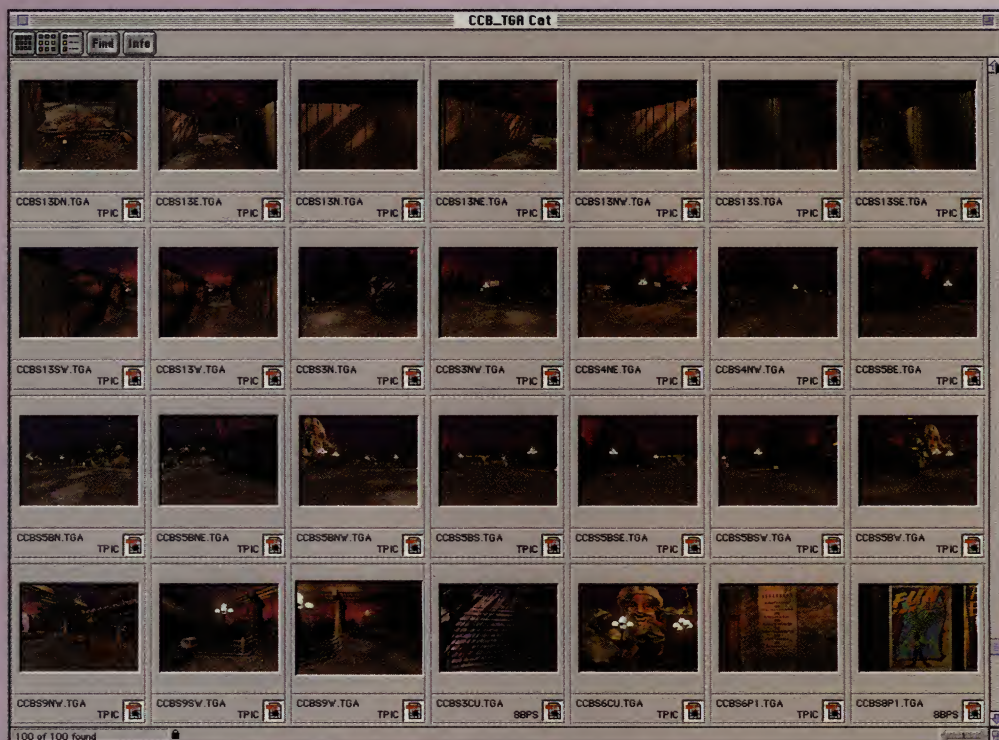
IP's online storage system includes a browser that allows the staff to search through assets. Shown here is a catalog of Concourse B still TGA-format images.

in, but also about it psychologically as it would appear to you.

Just looking at the maps it looks like these are specific navigation points.

Heinlein: Right. These are actual camera positions, and we knew we were going to make an eight-shot circle at each one. Because we were doing such a huge area, we had a huge palette problem in trying to make it appear realistic. I was not going to limit this to 256 [colors] for the entire game. Each individual shot has its own palette, but then we had the problem of trying to move the camera from position 1 to position 2, so we used a camera with a wider-angle lens so that each side of the next block over would overlap by roughly a third, so we would be pulling in these colors from the next frame over. There's a little bit of a jump as you move in a circle around a camera position, but I think it is pretty acceptable. It looks pretty good. And it's fast enough on a decent machine that I don't think anyone is really bothered by it.

Our camera lens was a 24. Then, of course, as we started developing it we could see that there were holes in our original concept, so all these little lines and red things [*points to tons of red pencil scribbles all over the maps*] were how many shots we added over the course of production. You

know you can sit here — and this is one of the things I tell my boss — "well, I can sit here and I can plan things till doomsday, but until you're in production and you can see the problem, you're never going to have everything nailed down tight."

In this concourse alone, we added whole camera positions, which are eight shots, probably like 20 to 30 more shots right here. That's just part of the fun. Actually, the red circles are when I got the shots back from the farm. I would have these laid out clean, and then I'd go through every shot and circle them off the list to make sure they were matching in color and to make sure everything was in its proper place. And finally, besides the excellent renderings the modelers did, we would always give it another shot in 2D just to blur some stuff out and pick some stuff up. A lot of my job was that, and signing off on these frames, making sure everything matched.

So we had six main sets on level one, which are the opening main gate, the four concourses, and the gazebo area. Then once you went through those, you got through a tunnel animation that took you to the interior office scenes, which was one of Tom's sets — Tom did all the interiors. We had the office set and the Oracle's chamber, so it was eight full areas. The bulk of the work was on the concourses. How many polygons were we running?

Miecznacowski: We had hundreds of objects. We were literally built down to the matchbook level where we had little pieces of origami that could animate. I would guess we had somewhere in the 100,000 plus range.

Heinlein: The concourses were like 700,000 polygons.

Miecznacowski: But there you're talking about acres. I'm talking about small rooms in the 100,000 plus range.

Heinlein: We also had a pretty good lighting system. We figured out that to make the glow off the lampposts in 3D Studio, we had to go in and anti-alias. Was it anti-aliasing?

Miecznacowski: We had to try a half dozen things to get the lampposts. The first thing we did was put maps on spheres, and it looked terrible. One of the other artists started experimenting with maps on planes. Those looked better. We tried IPASs, a lot of different things, and ultimately we ended up going back in and finessing each and every little detail. I think we ended up using maps on planes that rotated, and they added to the atmosphere. Shadow maps on lights, rather than the objects having shadows.

Heinlein: The one thing we wanted to avoid was ray tracing because of the rendering time.

Miecznacowski: Lighting was definitely one of the key things. That was where Doug helped me out. He really pushed my lighting. I would come in with a scene that I thought was great and then he would tell me, "No! Try this light here. More atmosphere for there." [*Laughs.*] Things like, use colored light to give something a sense of depth. For example, the Oracle room is a small chamber. We made sure we had cool and warm lights working against each other so there was real separation. Lighting was a huge endeavor.

Heinlein: Light being space, and colored light created emotional space. Just look at Hollywood movies to see some of that. We have a reel here of lighting effects that we put together from the popular media. It's just lighting effects that we show to freelancers. For me, it comes back to the painting background, just having art training. It's going back to when I was in Italy. I remember going in these dark churches where they have these boxes on the walls. You have to put money in them to turn on the lights. When I was about 17, I went to this one church. I put the money in and I turned around and there was this Caravaggio right in my face. It had this awesome lighting. I think it was *The Conversion of Saint Paul*, which was in Rome. That's what lighting is about. That's what makes stories happen.

It's interesting, too, because we think we might want to do some more twitchy kind of stuff, and they're asking us to back down a little bit in terms of technology. Are you willing to go to 33 colors? I think, "Yeah, probably for a year or so."

Twitch as in arcade-style games?

Heinlein: Twitch as in *Doom*. You know. It's really fast. It's really fun. It looks like crap. It drives me nuts. It's not acceptable. It's fine for now I guess, but it's not acceptable stuff.

You don't think hardware acceleration is going to help?

Heinlein: Probably, but to how many people?

Miecznacowski: The real hardware, GLiNT chips and stuff, that's a lot of money.

Yeah, but someone who wants to have the ultimate Doom experience might go out and spend a couple hundred bucks for it.

Miecznacowski: Also, the consumer has a different perception than we do. I do a lot of personal polling of game players in general, not necessarily just electronic gamers, and it's fascinating what their perceptions are. I was talking to a neophite 3D artist on the Net the other day and he was asking how come *Doom* could render so fast and other programs took hours to render a frame. They just don't understand.
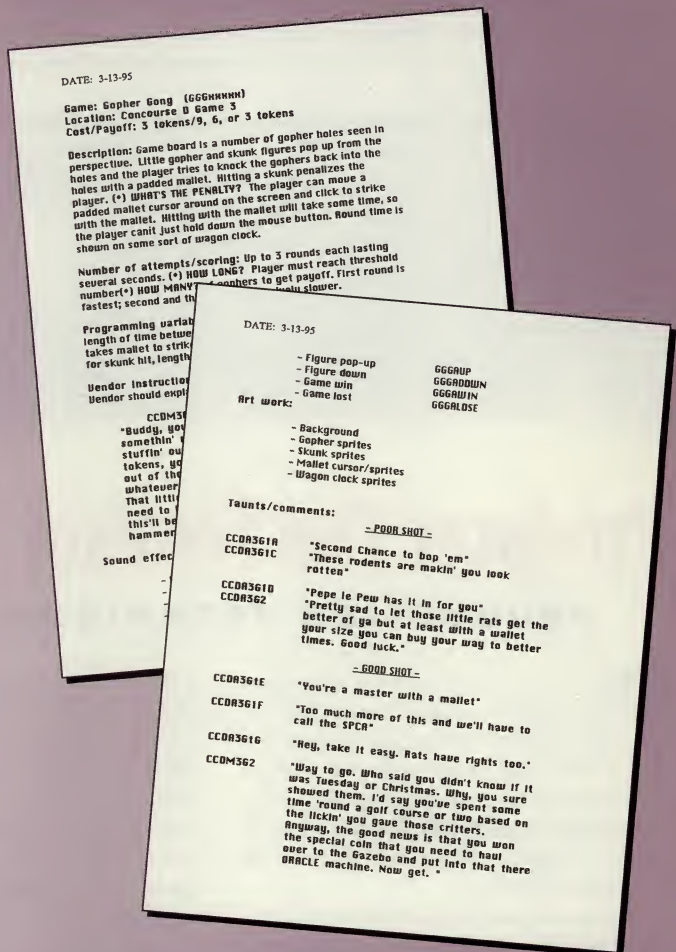
Examples of Howard Tullman's original script. Notice how he attached filenames to each asset from the very beginning. Also notice that the script was written not linearly but as a matrix.

Heinlein: I'm going to show you some more of this awful management stuff before we get into the art. This is how we track the nuts and bolts of how this stuff is put together. Movies, names, codes — when was it done, when was it due. We track everything through our asset management system. This is a model check-off so we have a whole ugly list, pages deep, of models, when stuff is final, who's doing it, breakout schedules. Really anal. Anal cool stuff. And I keep a wallboard on high-priority stuff.

Howard can tell you more about the characters and the scripting process, but I don't even have time to do full-color storyboards. We do try and storyboard as much as possible. The script is interesting in the way it brings out all the variations on each of the possible speeches. I work right off the script. It has the assets named already. Here [*points to a page of script*] we have the Water Gun vendor speech, we have the Water Gun vendor taunts, we have notes of changes, and we also have some of the logic developed already.

We don't really use email here, but we rely on voicemail heavily. At nine at night I have anywhere from 9 to 20 calls. Regularly, on a Saturday afternoon

was not determined at that point.

Heinlein: That's the nice thing about using 3D. If the light is in the wrong place, you move it. If you need another light you add it. That's why Hollywood uses the stuff. It's flexible.

There's equipment out there that lets you process everything in real time and that stuff costs something like $1,000 a minute to do. Unless you have a $6 million budget for your game, you're not going to do that. And then you're really talking about getting into Hollywood production. We try to be as production conscious and value conscious as we can. We definitely work hard to bring it off with the material that we are using. For whatever reason, our products haven't been in the marketplace as much as I would like to see them, and I know they're as good or better than anything else out there. I see video from other groups that is still sitting in boxes, it's one-eighth of a screen, and anyone can do that, that's no big deal. But you know trying to do what we're doing — full-screen video — is very difficult. And it's not the way we want it, either. We're still working hard with our research and in our alliances with other developers to get hot property first and be able to forge part of the avante-garde. We don't intend to be following anybody. We intend to set the standard higher, and from *Blown Away* and *Panic in the Park*, I think we've accomplished quite a bit.

Miecznacowski: Definitely. I think that's one of the first things you notice about our films. There is always room for improvement, but the acting, the characterization, and the dialog is definitely a cut above what's out there.

I will have 10 to 20 calls. That's really the way we communicate here. We have a voicemail nightly report for everyone on what's happening on the farm, what's happening in movie production, modeling — it keeps everyone up to date. Keeps everybody on the same page. We are so synergistic here, everybody is working so closely together, that we all kind of know where everyone is all the time. It's very tight.

Let's step back to the lighting issues. Were the 3D environments designed before the live video shoots were done?

Heinlein: We're working on such a quick schedule, we had the modeling blocked out for at least three quarters of it by the time they were ready to shoot video. There was a quarter that was just empty. There was nothing for them to go on so we had to build around what they shot.

Were you matching their lighting?

Miecznacowski: At that point we had given them cubes in spaces so that they would know where things were to get lens angles and so on, but lighting

Producer Tim Tennant

PRODUCTION TECHNIQUES, DEALING WITH HOLLYWOOD

Tim Tennant spent years as a production manager and assistant director for motion pictures and television commercials. Before bringing his production skills to Imagination Pilots, he ran Intergame, a joint venture with Digital Pictures, that produced games for Sega and Sony. They produced six games in six months.

"They said it couldn't be done, but we just hired the staffs we needed," Tennant explains. "The games were all based on movies. We hired large staffs and moved forward. What that experience taught me was that process is process.

It doesn't make any difference if you're making clothes or if you're making movies. You lay it out, identify your objectives, approach it as a team effort, hire the departments, and bring their expertise into it. That idea is foreign to the game guys, especially when they say, 'I can do it myself, I can do it myself, I can do it myself.' If you're going to do a number of games a year, you have to systemize the process."

The first thing Imagination Pilots did was develop their database, the asset management system. As Tennant puts it, "It lets the creative guys be free. The management system allows

Tennant brought Hollywood production techniques to Imagination Pilots. "The creative process is that much better because it isn't freeform . . . You learn to work within the time frame because that's all there is. Period."

them to focus their priorities. When you look at eight months on a wall, you can say, 'I know where I'm going, I know where I have to be, I know where I'm falling behind, and I know what's coming.'"

Putting that system in place was difficult at first. "We beat each other up for the first year," Tennant explains, "because everybody was set in their ways. All I could do was say, 'This is how we do it. This script of a movie is the same as this script of a game.' It's just information that an actor is going to read or the camera angle is going to follow. If it ain't on the page, it ain't on the stage. It's not going to magically appear."

Working with scripts written by Howard Tullman, however, poses an interesting challenge. "Howard is unique," Tennant says. "He can write a game from start to finish without building a flow chart. He'll write linearly, putting in the branching as he goes. And not only does he get real close to perfect the first time through, he includes all the logic as he goes."

That turns out to be backward, in terms of working out the production process, because Tennant has to reel Tullman's ideas back to fit real-world budgetary and time constraints. And while Tullman may have to give up an idea before it's executed, doing so in pre-production allows the team to apply their resources where they count and circumvent problems before they put the team over budget and/or behind

schedule. In the end, each person on the creative team knows what's expected of them on each day of production.

"The creative process is that much better because it isn't freeform. Everyone understands how what they do affects the other team members," Tennant tells us. "You learn to work within the time frame because that's all the time there is. Period. The first pre-production schedule for *Panic* had nine days of film shoot. We knew that would put us over budget by $250,000 to $300,000 just on the filming side." Tennant and Ken Berris put their heads together during the next few weeks and figured out how to get the schedule down to five days by omitting camera moves, matching actors to all their scenes so they wouldn't have to be on-set longer than they had to be, and so on.

To maximize the shooting schedule, Tennant arranged the storyboards in shooting sequence — the order in which they'd make most practical sense on-set as opposed to the order in which they'd appear in the game. No actor was on-set more than two days.

"There was a learning curve I had to go through," Tennant admits. "There are assets that they can't deal with until they're farther down the line, where I would have wanted to deal with

"Younger people aren't interested in people per se. They're interested in cartoon people . . . the vendors, the people who run the games in *Panic*, are cartoon people. They're abusive. They kid you. They joke," says Tullman.

them earlier. For instance, you can't post-score your music. You might think, 'Okay, the film is

locked, let's write the music for it.' But you can't do that without spending another four months trying to figure out how to do it. You have to do it as you go."

Composer Paul Libman (see page 58) understood that the approach required to score *Panic* matched the television commercial scoring model, not the feature film model. "The commercial people understand that things have to happen in exactly three seconds," says Tennant. "Commercial directors understand that they have 30 seconds to tell a story. From a creative standpoint, that fits the game model."

Tennant hired director Ken Berris to work on the games he produced for Sega and Sony as well as *Blown Away* and *Panic* for precisely that reason. "Ken was able to handle Howard, to structure Howard," explains Tennant. "Where Howard had written a paragraph of information in the script, Ken showed him how to cut it down to five words because that's all you need. And he handled Howard just like he would have handled a commercial client — 'yes Mr. Ford, I understand what you mean, but I want to show this red car this way,' and he tells a story visually.

"The trick is to hire pros, be they cameramen, music people, production designers, 3D artists, or programmers," he goes on. "There's a wealth of talent. You just have to know how to use them, control their budgets, and say 'I only need them for this amount of time.' That's the process. And it's just going to get better as the technology gets closer to television quality. It will just be telling stories on computers."

On the business and marketing side, Tennant says, "Anyone can develop a $60 million movie. It's a lot more challenging to take $1 mil-

TIM TENNANT on Budgets & Marketing

Panic cost just under $1 million to make. The initial budget was $923,000. We came in under that by about $10,000. I use a program called Movie Magic to track actual versus budgeted costs. To project real costs, I look at our components against time and money. One problem is that you can't get a definitive budget until a title is completely designed and scheduled out. You can't figure out the film elements until you've done that.

A day of filming is not a day of hard cost. You can't say five days is $200,000. They might be $500,000 or $150,000 depending on what you do. Each day has its own components. How much crew. How many actors.

The film portion of *Panic* was about $250,000. Erika cost us under $100,000, which included marketing and promotional time. She was worth it, but I don't think Warner used her to promote the game the way they could have. They wanted to wait and see how it did, but if I'm making a movie, I'm not going to wait for the third week to decide how I'm going to do my marketing.

This is a marketing-driven business right now. Unfortunately, there isn't an MTV out there selling games the way MTV sells records. But that will happen. When Intel makes a chip in a settop box that lets you sit there and play *Panic in the Park* on your television, then we've got a business. But that's going to take a while to happen.

The international market is exciting. I've got statistics that show in 1995, Japan drove from 3.5 million multimedia PCs in the home to 7 million by November. We think that will double this year. The growth in Europe and Asia will outsell the United States this year.

Germany is the biggest market in Europe followed by the U.K., France, Italy, and the Scandinavian countries. The issue is how many marketing dollars will we spend to support a title against a guarantee of pre-sale units. The fight is how much are you going to spend on my title before I give it to you. One side says there aren't a lot of quality titles on the foreign market, so they're buying whatever they can get. The other side says the international buyers won't pre-buy something.

We'll be doing all the localization at the same time as our production. We don't want to wait, so everything is done in sync. Soon the world will represent 70% of the marketplace and the U.S. will become just another market. If you don't realize that the international marketplace is your playing field, you're going to lose.

GOPHER GONG: From Start to Finish

The design process starts with the script but then creative director Doug Heinlein has to turn the idea into a concept that the modeling staff has to realize as 3D models, while interns animate the 2D sprites. You can see the script for Gopher Gong on page 53. Heinlein's concept sketches and the 3D models and 2D sprites they became are shown here.
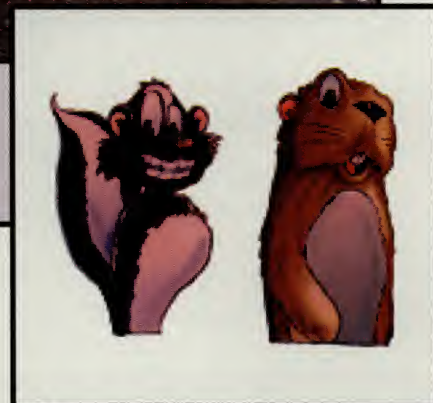
Heinlein's concept drawing of the game booth. The figure on the right gives the modelers an idea of relative height. The drawing style grew out of years of doing comps for ad agencies.

2D sprites of un-gonged and gonged critters.

The game itself. Above, the comp. At right, post 3D rendering. Premiere running on a Power Mac 8100 equipped with a Radius VideoVision Telecast was used to composite the sprites.

After all is said and designed, what happens next? IP programmers Mike Udrow and Phil Davis use a proprietary gaming engine that started life as a DOS tool and has since migrated to Windows. They are developing a new Windows-based object-oriented system. Their old engine's main limitation? Hotspots had to be rectangular.

PANIC IN THE PARK

Rock and Roll is taken from the kids game where you have to tilt a board to get BBs to lodge in holes. The BBs in this virtual version of the game obey the laws of nature, so test of skill is just as maddening as its real-world counterpart.

lion and do the same movie in six or seven months." So when it came to dealing with IP's lightning fast development cycle, Tennant says, "Warner Bros. had to go through a learning curve. They thought we were too fast from a marketing standpoint. They were saying they needed more lead time and we were saying, Why? Start when we start. You know what we're making. Here's the script. Get your box art done, start pre-selling, and seed the market with previews."

But Warner was nervous that consumer magazines were feeling burned by gamers who would announce titles and then not ship when they said they were going to ship. "Those rules

are going to change," says Tennant. "The guys who are going to survive in this business are the guys who say they're going to have it on Wednesday and they have it ready on Wednesday. Our cycles are very, very quick because we're a machine. We're like a studio."

Their unusual quickness has its downside, however. "We're in selling mode right now," Tennant explains. "When we show someone a treat-

ment, they don't want to respond overnight. They want to extend us out to 14 months. They want to decide by committee. No one wants to take responsibility, but the burn rate is too high. We don't have time to wait for them.

"You find that most of the people, especially from the entertainment studio side, haven't got a clue as to what to approve. They stall because they're trying to figure out what they're supposed

Paul Libman & Joel Fox

HOW TO SCORE AN INTERACTIVE ADVENTURE IN 3 WEEKS

For the musical portion of *Panic in the Park*, Imagination Pilots tapped the talents of Paul Libman and Joel Fox. Their primary business of scoring television and radio commercials made them perfect candidates for the job — the highly competitive world of jingle writing teaches you to work lightning fast and deliver projects on time and on budget.

From their studio, located a few short blocks from the IP offices, Libman and Fox were presented with two challenges. "The sheer volume of the work to be done," Libman explains, "and the tight deadline. We had three weeks to get almost half the game done so it could be previewed at a tradeshow."

From day one, Libman and Fox determined that they'd work directly in the digital domain, tossing out the thought of putting any of *Panic*'s music onto tape.

"Because the script is a matrix, we practically had to figure out a flow chart to keep track of everything in the game," says Libman. Even before seeing any of the actual film/game footage, Libman and Fox created a theme for every character. Then they previewed their initial ideas for Tullman and company, who gave them valuable feedback, which led to some midcourse corrections.

"Musically," says Libman, "we decided we needed to use live drums, guitar, sax, flute, clarinet. We recorded all of them into [Opcode's] StudioVision." StudioVision was also the sequencer of choice, while a Korg 05/RW and Roland JV80 supplied their General MIDI sound set.

Interestingly, Fox recorded the live players in mono, staying within their four-track StudioVision system for the recording. But they mixed the sequenced tracks and all the audio elements, including dialog, in an eight-track Digidesign ProTools

system. Each finished mix was put onto a Syquest cartridge and carried over to Imagination Pilots.

"I was sitting there composing as fast as I could," says Libman. "We did the first 155 cues in three weeks, including the live sessions with five different players. Our idea was to score each scene as if it were a scene in a movie. If we made any mistakes it was that in the beginning we were too specific with each scene because they weren't trimmed yet. We weren't getting final start and end points at first. But that changed, so it stopped being a problem." In

to say back to you. It's like when you're shopping a treatment and you give it to an executive. And you realize that they've it handed off to a reader. When you talk to them, they're reading off the cover notes. They haven't read the script."

But Tennant feels it's even worse on the game side. Film executives have no perspective on how to judge the progress of an interactive title. Warner execs apparently didn't even understand the overall premise of mystery in *Panic* or that the arcade games had little to do with the overall plot. As a result, Tennant says, "We have to fight tooth and nail to get them to understand that they're the money and the marketing; leave us alone to do what we do best. There's a phrase that I coined: defendable failure. How do they defend a game's failure if they haven't gone through the steps to protect their butts? And we say their defendable failure is that we have a proven track record, we deliver on time, we hired Erika. Stay out of the production process."

Part of the drive to unhinge themselves from Hollywood's partial paralysis is that the longer it takes to get to market, the greater the chance your technology will be yesterday's news before your title ships. Though on this point, Tullman and Tennant disagree. "My argument is that the story is king," says Tennant. "We're trying to convey emotion. I don't care about the puzzles. There has got to be a reason for you to play this game.

Panic's Quit menu gives players the option to save their game and/or change their mind. The sky seen here and throughout the park was created with KPT Bryce. The glowing light effect took the team a long time and a lot of experimentation to perfect.

Howard thinks it's technology, game play, story — in that order. I think it's all story."

The end result, luckily, is a give-and-take collaboration, which Tennant thinks is healthy. As he puts it, "I'm hopeful it will help us survive the shakeout that's coming."

---

fact, as they mixed the demos they were producing 50 to 60 customized pieces, which got thrown out because the scenes changed so much by the time they were finalized.

To keep themes from getting monotonous, Libman and Fox realized they'd have to do three or four different arrangements of each theme using different keys, different attitudes, different styles — so they could always find one cue that fit a particular scene. "We eventually realized that a library was the way to go," says Libman. "That was the key to make things happen fast. Customized cues became a last resort."

"One of the things you don't necessarily understand at first about music in an interactive game," Tullman points out, "is just how omnipresent the music can be. If all you do is loops, it starts to drive the players completely out of their minds."

Which was another reason Libman and Fox avoided using looping MIDI files and opted for WAV files. "If a character appeared 12 times," says Fox, "we'd do 12 variations." "That got a little tedious for us," adds Libman. "But that library concept helped."

Where they got creative and where things got fun was in the editing process. According to Libman, "We'd do crossfades, take solos from one scene and stick it over another piece, creating new compositions in the editing process."

On the technical level, in addition to recording in mono, lots of limiting (via a combination of analog limiters and L1, a TDM plug-in for ProTools) and equalization were used to maximize levels and brighten things up to compensate for the fact that they would never be heard in 16-bit, 44.1kHz high fidelity. Why mono? Besides not wanting to overtax their limited hard disk space, they knew they wouldn't have

the space to include full bandwidth stereo audio in the finished product.

Flattening out some of the actors' and actresses' voices proved another technical challenge that required the team to finesse certain pieces of dialog in Digidesign's Sound Designer. They were also careful to keep the dialog 6dB to 12dB higher than the music. "Some of the actors had gorgeous, booming voices that were difficult to record," says Fox. "The midrange component determines vocal intelligibility. It's not loudness. With one character, there was nothing the EQ could grab onto. That's when a combination of analog limiters and L1 came into play."

For down-sampling from 16-bit, 44.1kHz to 8-bit, 22kHz audio, they used Sound Designer and SoundEdit 16. But Libman says, "The difficult thing is there is no way to know in advance what something is going to sound like until you actually sample-rate convert it down. You know the high end is going to go away and that other things are going to jump out because they're in the right range. The bass actually came through pretty well, but I was surprised at how much of the high end went away in spite of the way we EQed the highs."

Sound effects were culled from an extensive collection of third-party libraries such as The Hollywood Edge, Twentieth Century Fox, and Hanna-Barbera libraries.

How did they mix music to video? Fox explains, "StudioVision didn't have a QuickTime window in it yet, so when a new Syquest would come in, we transferred all the QuickTime movies to 3/4" video. Then we striped timecode on the tape. The QuickTime movies included rough audio, which we kept on one track for reference. We'd drop in the good audio on a parallel track. With SMPTE on the tape, we could run StudioVision and ProTools in sync with the video."

PANIC IN THE PARK
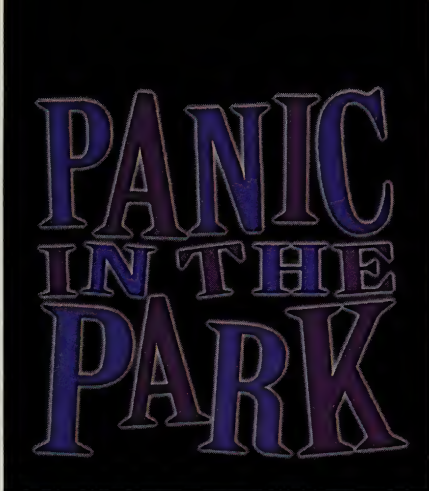
Contracts, SAG, The Director's Guild & Interactive Media

Three days prior to my visit to Imagination Pilots, I'd been hanging out on the set of a Trilobyte interactive film production featuring Hollywood talent [see *InterActivity*, Feb. '96]. On that set, I learned that the Screen Actor's Guild placed certain restrictions on what could be expected of an actor or actress. For example, their workday could not exceed 12 hours.

Considering that producer Tim Tennant is a member of the Director's Guild and Imagination Pilots has produced titles featuring Hollywood talent, I thought it might be interesting to get Tennant, Dan Di Caro, and Tullman's perspectives on working with SAG and the Director's Guild. Their responses:

Tim Tennant: Years ago, the day rate for a SAG actor was $2,500 a day. I was one of the producers who said, "This makes no economic sense. Why don't you just give us a contract that allows us to work with your talent and see how the industry develops." Now the day rate is $500 per day. The head of the interactive division of SAG was very helpful in leading the way for all the unions. He convinced them to learn what this industry is all about and paved the way for a contract that was very flexible. There's a day rate and a weekly rate for SAG actors now.

Howard writes games that have 15 to 25 characters in them. There's no time for error. We're doing 90 setups a day. Each actor has to come in and hit his marks, then move on. There's no room for non-pros. When you use a pro, you get so much more performance. The game comes alive.

Dan Di Caro: The director's side of the equation is separate from the SAG side. What the Director's Guild is trying to do is have you sign their television contract, then add a waiver that says you can work within certain parameters, get points, and so on.

We've refused to do that until they orchestrate an interactive contract. Their objective is to get a bunch of companies to grandfather those contracts. Different companies have different views on this. The folks at Hyperbole are very union oriented, but I don't know what that translates to when you have Greg Roach sitting there being writer, director, and a million other things.

These are serious issues because until the economics of this business improve substantially, having 30% to 40% fringes added on the already pretty healthy day rates makes live action production very costly.

One of the questions I have is that if some of these guys build inhouse facilities, they could do it all within the four walls of their business. So the Hollywood mentality has got to adapt to a new genre for a period of time.

Howard Tullman: It's true the video game industry is large, but the video game industry is not the CD-ROM business. The thing that these Hollywood guys don't understand is that CD-ROM is a very niche business. You can talk all you want about the installed base of multimedia machines, but until it's so gigantic that one or two percent represents hundreds of thousands of units, the math isn't going to be there.

Another thing Hollywood doesn't understand is that Sega and Nintendo have an absolute stranglehold on the technology, so you end up paying 30% of the cost of the product to them as a licensing fee. That's not true in the CD-ROM business, so that whole premium profitability of the game industry has gone away.

People also don't understand that there are a lot of alternatives to buying a computer and a CD-ROM. Books, movies, whatever. I don't think the pool of entertainment dollars is infinitely expandable. But the biggest bite in the butt is that the principal consumer of consumer media is women. And here we are with a complete demographic mismatch. We're trying to figure out how to sell games to guys.

People say the market is in the billions, but to put things in perspective, guys like Michael Eisner preside over businesses that might do $50 billion. That sounds like a big number, but Wal-Mart did $85 billion dollars in total revenue last year.

CFO Dan Di Caro

ON MANAGING FINANCIAL STUFF

My background is years of financial management. Howard Tullman and I have been together for about eight years. Our arrangement with Time Warner is that they are a part owner of this venture. They provide equity capital and fund development of our products. We still have to submit proposals for approval to get financing.

That has worked out well so far, but we're in the process of trying to go beyond that arrangement to do cinematic titles with other publishers. There are a lot of reasons for that. We think it just makes sense to broaden out horizons, because their focus isn't high-profile film-based titles. The *Where's Waldo* series is more their style, and they're delighted with that series.

The riskiest part of our operation from a financial standpoint is the film production costs. If it had rained on one of those days — if something had prevented shooting — I don't know what we could have done. It was that tight. By the time we're ready to start shooting, it's like you've flipped a switch and there's no turning back.

The other difficult thing is balancing star quality versus the cost. So far, the best thing that Erika's presence has done for *Panic* is get us a link on the Playboy Internet site (http://www.playboy.com/). Once that happened, traffic went through the roof. We were getting up to 60,000 people a day looking at our site. One of the nice things about having a visible actress is that it will also translate into domestic and international sales.

It's too early to share sales figures with you. It takes a minimum of 60 days from the end of the month to get sales figures from our distributors in the U.S. and 45 days from the end of the quarter from our international distributors. The time lag is tough. Time Warner has a distribution model that could modernize the way sales are reported. I'm hoping they can lead the way for the industry.
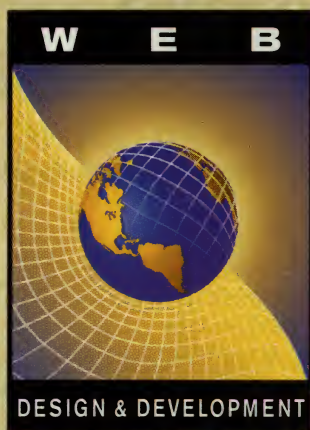
A DESKTOP VIDEO PIONEER RETOOLS

RADIUS VIDEOVISION TELECAST AUDIO/VIDEO CAPTURE/PLAYBACK FOR THE MAC

BY CHRIS MEYER

Today it's nearly taken for granted, but the ability to integrate video with multimedia content and record the results directly to videotape is fairly recent. Formerly this process required either file transfers to and from external hardware, expensive special-purpose cards, or putting up with lousy image quality.

Now several full-motion video cards are available to do the job. They get more notice for putting nonlinear video editing on the desktop, but these cards are also useful tools for creating multimedia and outputting to tape short video segments such as titles, show openers, and special effects. These are the applications we'll focus on here.

Radius VideoVision is one of the longest-lived products of the genre. The latest incarnation, VideoVision Telecast, works with a NuBus-capable Macintosh to capture and output broadcast quality video. Rather than review Telecast in isolation, let's consider it in light of the evolution of the VideoVision line and desktop video in general, including issues, problems, and solutions.

The Dawn of QuickTime

When multimedia was just starting to gain momentum, QuickTime and Video For Windows hit — and suddenly it became obvious that even your pet parakeet would want to embed video in spreadsheets, word processing documents, and business presentations. Apple was busy developing a chip known as Touchstone that would help enable such applications. Eventually they sold the technology to a third party, Radius.

Radius released the original VideoVision in mid-1992 for a list price of $2,000 ($2,500 bundled with software). It consisted of a NuBus card and an external breakout bar for video interfacing. The card itself sported a Mac-standard RGB video output that could drive a 640x480 monitor at 24-bit color depth or larger monitors at less depth. This output could be configured for either NTSC frame rate (29.97fps) or PAL rate (25fps), with interlaced fields per frame (two images for each frame, with even and odd scan lines drawn separately) rather than the progressive scan format used by computer monitors (a single image

made up of lines drawn sequentially).

The card's local connector could output only RGB format, which is not compatible with most video equipment. It was the break-

Figure 1a.

Figure 1b.

out bar that featured the real video and audio connectors. Through this bar, the system could output both NTSC and PAL interlaced video in composite and Y/C formats (see

Figure 1. The results of (a) zooming from 320x240 to 640x480 by repeating neighboring pixels, as most video cards do, and (b) VideoVision's interpolated zoom, which usually results in a more palatable image, defocused rather than blocky.

"Video Interconnection Formats," p. 64). It could also capture video through two pairs of composite and Y/C inputs.

The breakout bar included stereo Audio Ins, although they were restricted to then-standard 8-bit resolution and a 22kHz sample rate. It also included a pair of Audio Outs — presumably to feed the video deck to which you were recording the computer's output — but they delivered a mix of the signals sent to the Audio Ins and Mix Ins, not audio from the Mac. To do this, you had to run a cable from the Mac's own audio outputs to the bar's Mix In connectors.

Beyond combining video inputs and outputs in one system, a few other features set the VideoVision apart. It was one of the first systems designed expressly to capture and play QuickTime movies — although in those early days, everyone's frame rate was low and their frame size small. More significant was *convolution*, the ability to blur the NTSC or PAL output slightly so thin lines (such as those that make up the outlines of most computer windows) wouldn't flicker excessively when they were displayed on interlaced video monitors or recorded to videotape. This made Video-Vision well suited for outputting computer-based presentations. It also made any monitor connected to it usable as a second computer display, even an interlaced video monitor.

Radius promised expansion capabilities, including an extra serial port on the breakout bar and two expansion connectors on the card itself. But, of course, savvy walkers of the desktop video tightrope didn't hold their breath.

Full-Motion Video Arrives

One year later, Radius followed through with a $2,000 add-on, the Studio daughtercard. The new card sported a JPEG compression chip that enabled full-screen, full-motion (25fps or 29.97fps) video to be captured and played back in conjunction with any Quick-Time-compatible application. It was not the first card to offer this capability, but it did have some features that separated it from the pack.

Most important, VideoVision Studio (VVS) was the first commonly available card for the Mac to support both fields of an interlaced video signal properly. Many of the early "full-screen, full-motion" video cards actually weren't able to capture two fields per frame running at 59.94fps. Instead, they ignored one field and duplicated lines to fill out the image. This required less data but cut both horizontal and motion resolution in half. Some cards took another route and captured both fields but threw away every other horizontal pixel and duplicated its neighbor to fill out the image. The latter approach resulted in a noticeable loss of vertical resolution.
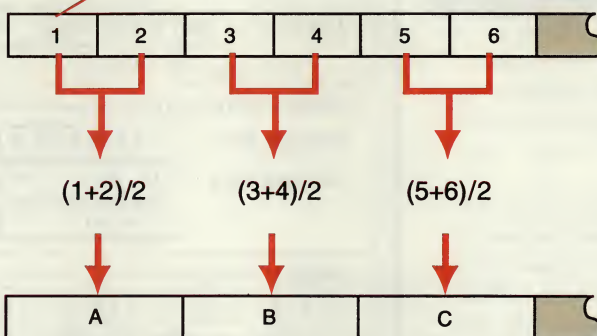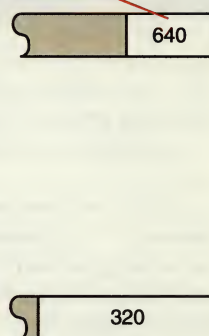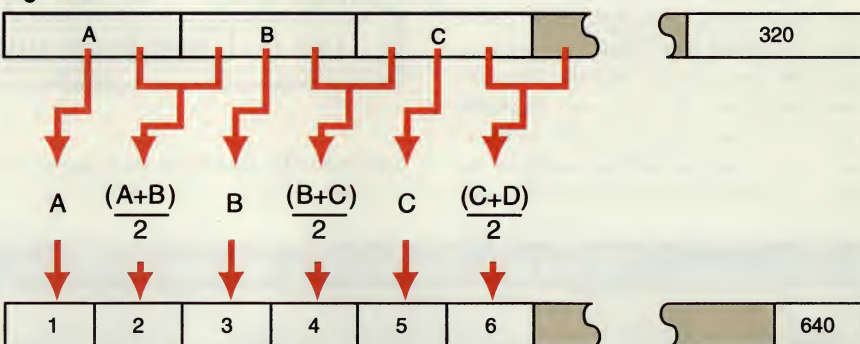
Figure 2. With Horizontal Interpolation enabled during capture (a), adjacent pixels are averaged and stored 320 rather than 640 per scan line. During playback (b), the 320 pixels are spread out over 640 in a way that minimizes loss of detail.

Cards resorted to these tactics because they weren't fast enough. That is, the degree of data compression necessary for them to move two video fields at full-screen, full-motion resolution incurred unacceptable degradation of the image (see "Crunch Time," p. 65). VVS could do the job, and not only that, it provided better performance in a number of more demanding situations.

For example, many cards could play a less data-intensive 320x240 movie and zoom it to fill the entire screen. Unfortunately, they usually did this by duplicating every vertical pixel and horizontal line, resulting in a very blocky image. By contrast, VVS offered hardware-accelerated *interpolated zoom*. This technique made the image look defocused rather than blocky (see Figs. 1a and 1b) — a better compromise and a rare feature even today. (Interpolated zoom shouldn't be confused with convolution. VVS automatically turns off convolution when playing back full-motion video to avoid blurring it unnecessarily.) VVS also offers a mode in which only 320 pixels are stored per horizontal scan line, but

instead of dropping every other one during input and duplicating what's left during output, it combines and re-interpolates them in a way that results in less-offensive artifacts. The difference is illustrated in Figs. 2a and 2b.

Based on price, capabilities, and QuickTime compatibility, Radius VVS was the first system to prove the potential of desktop video on a budget. It was used in both industrial and actual broadcast applications. But there were still some shortcomings.

For starters, having only composite and Y/C video connections limited the maximum quality possible. Even though a large number of studios still use only composite-level interconnects, component video is required for more critical applications. Radius supplied a partial fix in the form of a genlock/RGB cable kit. With the RGB cable, the RGB NTSC output from the main card could be transcoded via one of the many third-party boxes available into whatever format you could afford. Still, this provided only component video output; there's no RGB input on the VideoVision card. The genlock cable turned the extra serial port

on the breakout bar into a sync input, enabling VVS to be slaved to a studio's blackburst or house sync generator. In a professional studio, virtually all gear is *genlocked* to a master sync generator.

A bigger issue involved missing pixels. Professional NTSC video equipment specifies 486 active horizontal scan lines as opposed to 480 with VVS and most other cards. The official CCIR-601 digital video specification also calls out 720 vertical pixels, not 640. Generally speaking, desktop video mavens must learn to live within these constraints, but they're a compromise.

What most users couldn't live without was high-quality audio. Simply put, 8-bit, 22kHz audio restricted the system to low-end applications. Even more problematic: The audio wasn't synchronized to the video, inevitably resulting in loss of sync during the course of longer programs. You could use a separate audio card or the built-in 16-bit 44.1kHz audio with newer Mac models, but there was no way to synchronize it (see "Two Ships in the Night," p. 66). Indeed, audio/video sync is one area where both software and hardware con-
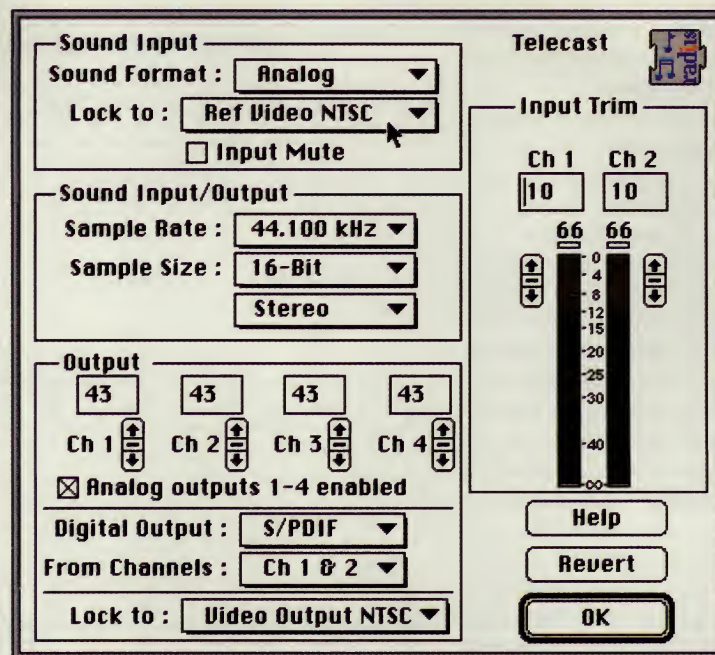
Figure 3. VideoVision Telecast's Sound Input control panel, where synchronization, level trimming, and audio formatting chores are performed.

VIDEO INTERCONNECTION FORMATS

Video comes in several formats. The *composite* format is most common. It's the signal sent by your VCR or camcorder, and it's supported on virtually all television monitors and video decks of all formats. It's called composite because the entire video stream is encoded as one signal and therefore travels through one wire. The result doesn't look great, but it's what most of us are accustomed to viewing on TV. Composite video is associated with phono (a.k.a. RCA) connectors on consumer and prosumer equipment and BNC connectors on professional equipment.

Y/C format encodes video in two signals, brightness and color (*luminance* and *chrominance*). It looks considerably better. It's the format used by S-VHS and Hi-8 decks and by some higher-end equipment as well. The usual connector for Y/C is the same as the one found on Macintosh ADB cables. On rare occasion, you find Y/C delivered by a pair of BNC connectors.

Higher-end equipment usually supports *component* video. Most often this is delivered via a set of three BNC connectors; a fourth may be used to carry sync. Unlike the RGB video found in computers, in which the signal is made up of red, green, and blue components, the three cables of component video carry a signal known as *YUV* made up of luminance plus two chrominance signals that are different combinations of the red, green, and blue. The result is cleaner, sharper, and more accurate than Y/C or composite. Virtually all Beta SP tape decks — perhaps the most commonly used tape decks in professional video today — offer component video in addition to composite and occasionally Y/C.

Bear in mind that component video is more expensive to implement. It isn't common to see a studio in which every piece of equipment is connected using this format.

It is important not to confuse video signal formats and tape formats. A Y/C signal sent from a VHS deck to a monitor will look better than a composite signal — the tape format is identical, but the higher-grade signal format makes a difference. Likewise, you can downgrade even the best video equipment by connecting it with composite signals. I've encountered at least one person who insists on using composite video instead of Y/C strictly because composite was delivered via a professional connector and Y/C by a consumer connector. His loss.

All the above are analog formats. High-end digital video equipment, including D1 or Digital Betacam VTRs, pass both audio and video in digital form along a single serial cable. The format might be component color (as is the case with D1 decks) or composite (as is the case with D2 decks), but digital composite video with analog composite video are not comparable. Similarly, the data on the tape itself might be compressed, as Digital Betacam is. However, the digital video that passes between decks is uncompressed.

Prosumer digital video equipment is starting to appear that carries video on a P1394 or Firewire serial cable. This is different from the format used by D1 decks and their ilk. What's notable about the new format is that the digital video signal is compressed both on tape and traveling down the cable. Also, it doesn't permit the color resolution of higher-end formats. (For the techies out there, it's 4:1:1 YUV for NTSC as opposed to 4:2:2). Nonetheless, early indications are that this new digital format looks as good as some of the best analog formats. Being digital, it offers the bonus of being computer-friendly from the start. ✸

tinue to fall short throughout the desktop video market.

High Quality Audio & Video with Telecast

In early 1994, Radius displayed another add-on NuBus card called Telecast, intended to occupy the VideoVision card's second expansion connector. The card comes with a two-space rackmount box that replaces the original breakout bar. At the time, it was a $12,000 add-on to VVS (then $4,500), making the price of a full-blown VideoVision Telecast (VVT) system more than $15,000. It didn't start shipping in quantity until 1995.

The high price and some nagging bugs — not to mention increasing competition — prevented sales from taking off. Today, Telecast's list price has been slashed to $6,399 (VVS is listed at $2,999). The bugs have been ironed out, making the system a serious tool for video production all the way from CD-ROM to broadcast.

Telecast adds many important capabilities to VVS, primary among them are component video connections and synchronized CD-quality audio. Component video is used by higher-end video equipment such as Beta SP VTRs (see "Video Interconnection Formats," p. 64). This results in significantly higher video quality both coming in and going out. Even video captured for low data rate CD-ROM delivery benefits from this. Higher-quality video compresses more efficiently, resulting in better-looking movies without an increased data rate (see Frames of Reference, *InterActivity*, July/August '95).

The new audio capabilities include two channels in and four channels out of pro-format analog audio plus two channels each way of pro- or consumer-format digital audio (see "Audio Formats of the Rich and Famous," p. 68). A headphone jack adorns the front of the rack box, and a monophonic mix output on the back can be routed to inputs available on many video monitors — a nice touch.

Sample rates include PC- and Mac-standard 22.050kHz and 22.254kHz as well as 32kHz (featured in some European broadcast systems), 44.1kHz (the rate for audio CDs), and 48kHz (found in DAT and digital video decks). The system supplies pull-up and pull-down variations of the last three, useful when you're interfacing with systems that send 30fps rather than the actual NTSC video rate of 29.97fps. Figure 3 shows these options along with the detailed signal level meters and gain adjustments available in VVT's SOUND INPUT control panel.

Separate options are available for syncing the sample clocks of the digital audio input and output. Input can be synchronized to an internal crystal, incoming digital audio, incoming video, or the genlock input in either NTSC

CRUNCH TIME

It requires 900kB to represent a 640x480 24-bit RGB video frame in the digital domain without throwing away data (otherwise known as compressing). This means a computer and/or hard disk playing uncompressed digital video at 29.97fps must accomplish nearly 27MB per second of sustained throughput. If you keep the data in its native YUV format rather than converting to RGB, the figure drops to 18MB per second (that is, 600kB per 640x480 frame of 4:2:2-format 8-bit luminance YUV video). Supporting the CCIR-601 format of 720x486 pixels and 10-bit luminance YUV raises it back to approximately 25.5MB per second. Any of these numbers is pretty high for today's computers.

This is why video is usually compressed for the desktop. Most Mac and PC video cards use the JPEG compression algorithm, so at least we can compare apples to apples when it comes to their respective compression ratios. (Nonetheless, there are different varieties of apples. Different cards use a number of tricks to maximize the technology, to greater or lesser effect.)

Early in the full-motion game it was the cards, not the computers and hard drives, that posed the greatest limitation. For example, the SuperMac Digital Film card pumped out close to 1MB per second. This works out to a compression ratio of nearly 27:1 (based on a 640x480 RGB source), resulting in unacceptable JPEG compression artifacts. Some cards dropped fields (cutting both horizontal and temporal resolution in half) or cut vertical resolution in half, halving the amount of data to be compressed. These techniques might be considered better if what's bothering you is JPEG artifacts, but reduced resolution also degrades the image. Scaling down the image resolution to 320x240 cuts the data rate (or amount of JPEG compression required) by a factor of four, but the image can look even worse — especially if interpolated zoom (a feature of the VideoVision line) isn't available.

Radius VVS itself doesn't pose a meaningful limitation in data transfer rate, so the issue shifts to hard disk speed. Raw disk speed is not an accurate measure because the file system and QuickTime impose significant overhead. Fortunately, VVS's FIND command measures a system's overall performance, and its adaptive compression scheme rides the compression ratio to use the throughput available as efficiently as possible.

According to the FIND command, the Macs of a few years ago were capable of transferring 1.5MB to 2MB per second from a single hard disk. The disk arrays of the same era could double that number. This made it possible to work with compression ratios in the area of 15:1 to 8:1 — borderline acceptable for industrial and broadcast applications. For short segments in which quality was important, I used a RAM disk to achieve data transfer rates of up to 6MB per second from a Quadra 950.

Today, those numbers have nearly doubled again. The FIND command says I'm getting 3MB per second from a speedy single drive, 4MB to more than 5MB per second from a disk array, and 9MB per second with a RAM disk. At these rates, JPEG artifacts have become harder to see. Now the primary limiting factor is the quality of the video interconnections and the chips on the card itself.

The newer generation of boards have better chips and component-format video interconnects, both of which increase the subjective quality of the output. These features also reduce the amount of noise at the video input, which gives codecs trouble, further increasing the efficiency of compression. As a result, captured images look better at the same sustained transfer rate (or, conversely, a lower transfer rate can be used without degrading image quality further).

How high a transfer rate you need depends on the quality of your entire video subsystem as well as what the original video looks like. For example, busy video can be hard to compress but it can also hide artifacts. On the other hand, something as simple as white text against a black background compresses more efficiently but also exposes artifacts, even at the highest transfer rates.

As noted in the review, at CyberMotion we make do with transfer rates of 3.5MB to 4MB per second using component video interconnections between the VTR and VideoVision Telecast (with a Power Mac 8100/110 and disk array). I'd like to increase that to 6MB per second (150kB per frame at 640x480) before I stop worrying. And even as good as that may look, it pales in comparison with uncompressed 10-bit CCIR-601 video from the latest Abekas hard disk recorders. ❦

TELECAST

or PAL modes. (The Telecast I tested strongly preferred blackburst sync rather than composite sync, which created trouble with the audio input. Most sync generators, including the Horita BSG-50 that I used, provide both.) Telecast has a built-in sample-rate converter in case you want to digitize audio at a rate different from that of incoming digital audio. If you use this facility, input can be set to synchronize to the internal crystal or the video output. Syncing both input and output to video reduces the chance that things will fall out of sync.

In addition, the system provides SMPTE input and output connections that support both LTC and VITC timecode formats — another advantage over most competing systems. This comes in handy when using VVT for video editing or as a digital VTR, when maintaining sync with various decks is necessary.

(We don't need these capabilities at Cyber-Motion. Except for quickie rough-draft proofs, we only record to and play back from Beta SP, which embeds timecode in its serial control signal, so SMPTE read/write is built into our VTR control. For capture and playback, we use Adobe Premiere and Pipeline Digital's ProVTR and Studio Recorder software and cable — all formerly bundled with Telecast; Studio Recorder is now offered as an option — all of which derive time from the Beta's serial cable. Incidentally, both Beta SP and Telecast have

component video I/O, which makes for an ideal marriage.)

Of course, there are still some drawbacks, despite the price reduction. For a start, Telecast is happiest in a limited number of Mac models. Radius lists the Quadra 840AV, Quadra 950, Power Mac 8100 or 8100AV, and Workgroup Server 9150. Indeed, both VVS and VVT had

8100/110 with a disk array, compared with more than 5MB per second in an antiquated 840AV.

Furthermore, the newest Mac models accommodate PCI rather than NuBus cards. As of this writing, Radius has announced a PCI version of VVS, expected to ship by the time you read this, that features synchronized 16-
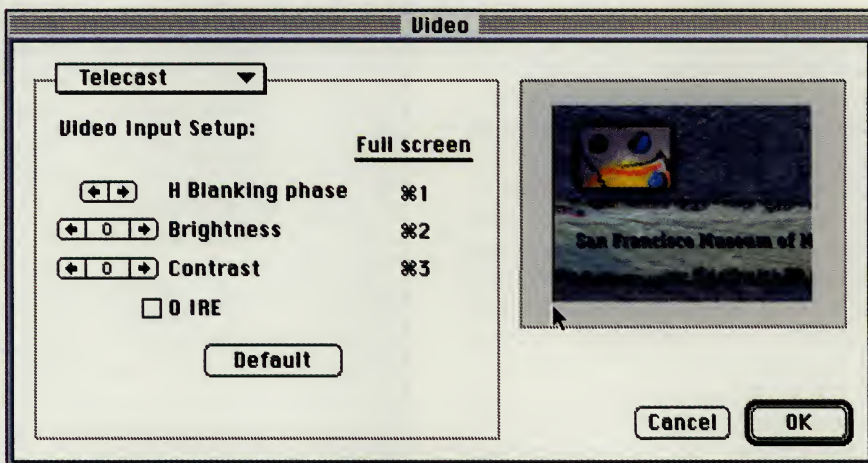
Figure 4. Trimming VideoVision Telecast's horizontal timing to eliminate vertical black bars along the left or right side of captured video.

trouble with the Power Mac models when they first appeared. Even with a hardware update for the VideoVision card plus software and firmware revisions, maximum data transfer rates on Power Macs aren't as high as I would expect, judging from competing products. I got around 3.5MB to 4MB per second in an

bit audio, improved data rates and image quality, and file compatibility with NuBus-based systems. A PCI-based successor to VVT featuring component video is said to be in the works, but no further details are available.

The issue of supporting a 640x480 pixel display rather than the professional CCIR-601

T W O   S H I P S   I N   T H E   N I G H T

In any video system, audio and video need to be synchronized so the image and sound line up properly — a simple enough fact. When you're recording video to tape, it's relatively easy: The two signals are recorded together on the same physical tape segment. In fact, there's no way to separate them.

It isn't so easy inside a computer, where audio and video are usually processed separately and stored on different parts of the hard drive. You can tell them to start and stop at the same time, but you have to take extra steps if you want to make sure they don't wander off in between. Early desktop video systems such as Radius VideoVision Studio didn't take these steps, and the twain did indeed wander. Furthermore, most computer operating systems don't allow the computer's own internal audio to be controlled by external hardware, so any video system that relies on a computer's built-in audio also is prone to wandering.

Various forms of trickery can be applied to do the job by brute force. Some video systems drop or repeat either audio samples or video frames

to keep the two lined up. This method works but introduces undesirable artifacts. Consequently, in one of our computers we use a separate audio card that can synchronize to a video signal (i.e., the one being output by my video card).

The best solution is to use a single card that hard-synchronizes audio and video. Several newer systems, including VideoVision Telecast, do this. Nonetheless, it can't be taken for granted that all systems claiming audio/video sync use the proper gear box between the two or that the clutch between them never slips. Furthermore, software issues exist that can trip up even the best hardware.

A technical standard entitled SMPTE 272M describes the proper synchronization between audio and video in a digital video system. We'll take a good look at it in next month's Frames of Reference column. Meanwhile, grill the hardware (and software) manufacturers you deal with to determine whether, when they claim synchronized or genlocked audio and video, they follow the 272M spec. If they don't, let 'em have it.

720x486 format remains. It's livable, but it's a compromise, and it makes combining Telecast images with those from an uncompressed digital video recorder such as those made by Abekas, which do follow CCIR-601, more difficult. (For various workarounds, see Frames of Reference, *InterActivity*, May/June 1995.) Furthermore, VVT doesn't support on-card hardware acceleration for video transitions in any editing application, which some of its competitors do. At CyberMotion, we tend to use software that's not accelerated by any video boards, so this isn't an issue for us. It might be if you use Adobe Premiere (for which acceleration is provided by some video cards) or some other editing program dedicated to a specific card.

Life with Telecast

We already owned a pair of VVS systems before Telecast became available. This made adding VVT to a third computer (for us, a Power Mac 8100/110) an easy decision, since we're familiar with the basic system and it's data-compatible with our other workstations and archived material.

Despite some horror stories I heard early on, installing Telecast was close to an out-of-box experience. We did need to get an older VVS updated to work with the PowerPC NuBus slots, but that was relatively painless. We teamed the 8100/110 and VVT with a high-speed disk array and Beta SP VTR, so it also functions as our main transfer station.

In some professional video environments, the machine room with all the gear is a good distance away from the actual work area. This reduces noise and makes it easier to share gear between adjacent studios. You might also want to move the computer and disk drive, with their noisy fans, away from your workstation. Taking this into account, the extra-thick cable that connects the Telecast box to the VVS and Telecast cards in the computer is a whopping 15' long, and 15' extensions for the ADB and ProVTR serial cables also are included. In our case, the workstation, computer, and equipment rack are within a 5' radius of each other, so the bigger problem was finding a place to coil the python-like cable.

The Telecast rack box can be mounted with either the connectors or panel indicators facing out. I prefer the latter. The indicator panel is adorned with status LEDs that tell you what audio and video formats currently are selected plus rudimentary signal level meters. The box I used exhibited the odd habit of lighting the audio's SYNC TO VIDEO indicator only after I had captured video or fiddled with the control panels. However, all the audio settings including this one were retained properly between power-ups. When operating, the Telecast box has a whisper-quiet fan, and when the computer is off, a single STANDBY indicator remains lit.

The manual suggests you leave an extra rack space above and below the box for adequate cooling. I installed it in an open-frame rack with shallow equipment mounted above and below it and encountered no problems.

As helpful as the box's LEDs can be, the various software control panels and Premiere plug-ins provide considerably more information and control over VVT. An audio output control panel joins the audio input control panel shown in Fig. 3. When capturing video within Premiere, VVT enables you to set the video input as well as sync mode (genlocking to an external master or following the video signal itself). You can also do other useful adjustments such as horizontal offset to remove black stripes along the left or right edges of a captured image (see Fig. 4). Retained from VVS are convolution (again, turned off during playback to avoid blurring the image), smooth zooming, and an option to guess your hard disk's capture rate and clamp capture and rendering at that rate to avoid dropped frames.

The Score

There are special considerations when you're selecting a desktop system for functions other than editing. No current product answers all my wishes. Here's how Radius Video-Vision Telecast scores on several criteria.

First, the good news:

▶ It's QuickTime compatible. This makes VVT compatible with a wide range of Mac software. I use many programs and rely on QuickTime to move files among them.

▶ It's scalable. I can install a full-blown VVT in my main video transfer station and compatible, less expensive VVS systems elsewhere. It can output passable video to a normal single disk drive, especially with the smoothly interpolated hardware zoom, and

it can output even better video with an array. Better yet, Radius offers a program called SoftStudio, a software-only codec that reads and writes VVS/VVT files without requiring compatible hardware on the rendering station. It can be downloaded free of charge from Radius' Web site (http://www. radius.com). This will make life easier within our walls and help us integrate freelancers as the need arises.

▶ It synchronizes audio and video. This capability was sorely lacking in VVS. The system offers a helpful range of options for sync source, plus sample-rate conversion and pull-up and pull-down rates.

▶ It supports another video monitor. Not all cards allow you to use their NTSC or PAL output to connect a fully functional monitor. We always have our cards connected to an NTSC monitor as a second screen to check colors, as a workstage, or to just hold extra menu items and palettes. Furthermore, convolution keeps

it from flickering. Most NTSC cards lack antiflicker circuitry or double up horizontal scan lines when they're not playing back video. ·

▶ The quality is good enough for broadcast applications. The audio and video connections are pro-format. The video looks good. The audio sounds fine. Everything syncs nicely. An added bonus is digital audio interconnections for DAT recorders, which most competing cards don't provide.

That said, I do have a nit to pick.

▶ There's no support for image resolution of 720x486 pixels, as described in the CCIR-601 digital video spec. This is the standard for professional video recorders as well as Abekas-format video frames.

At CyberMotion, we've been using it daily and we're happy with it. As one of the best-established desktop video systems on the market, VideoVision is likely to be resident in other studios with which you may need to trade files. Also, the recent price cuts make it a very attractive alternative to its competitors.

On the other hand, most of those competitors are a good deal younger, and this gives me pause. While Radius gives good, honest support and has upgraded the VideoVision line continually over the years, there are indications that the system may be getting long in the tooth. It's fussy about what computer you use it with. Details on the PCI upgrade path for VVT aren't yet available. And several key engineers who created the original VideoVision are no longer with the company. I've enjoyed the ride over the past four years, but I wonder how much further Radius can take it before they need to start fresh.

But these concerns are intangible. VideoVision Telecast is here now, and it delivers. If you're looking for a good-quality video capture and playback station and are happy to stay with NuBus-based Macintoshes for this function, this system is one of your best options. &

AUDIO FORMATS OF THE RICH AND FAMOUS

Most consumer hi-fi and video systems use a format of analog audio known as line level, unbalanced. The unbalanced part means two wires are used per audio channel, which carry the signal and the ground reference, often via phono (a.k.a. RCA) connectors. Line level means the voltage reference level used to indicate a nominal volume level is fairly low. In audio circles, you'll hear the term -10dB, which is identical with line level.

Professional audio and video equipment uses a format of analog audio known as +4, balanced. The balanced part means that three wires are used instead of two: a positive version of the signal, an inverted negative version of the signal, and a ground reference. The two phase-inverted versions of the signal enable balanced systems to reject hum and noise that accrues within the cable itself. The connector format used most often is XLR (a.k.a Canon). The +4 part refers to the higher-voltage reference level used to indicate a nominal volume, which also helps overcome noise.

VideoVision Studio, consumer VCRs, and most computers use line level, unbalanced audio connections. VideoVision Telecast and professional VTRs support +4, balanced audio connections. What if you need to use line level, unbalanced equipment in a VVT-based system (or, for that matter, professional VTRs in a VVS-based system)? At the very least, you need adapter cables to convert between balanced and unbalanced lines. VVT's audio control panels enable you to adjust signal level between +4dB and -10dB. Some VTRs, such as the UVW-1800, enable you to switch levels on the back as well. However, in some situations you'll need a special device known as a match box to match -10dB and +4dB signals — otherwise, distortion and/or excessive noise can result. Several manufacturers make line level match boxes. Ebtech's rackmount unit that converts eight channels of audio between these formats and levels (around $300) is a worthwhile investment.

Digital audio also comes in consumer and pro formats known as S/PDIF and AES/EBU. Both use the same raw format for the audio data itself but implement it in different ways. Typically, S/PDIF is associated with phono connectors, uses a lower voltage reference for its digital ones compared to its zeros, and uses a slightly different pattern of ones and zeros than AES/EBU for non-audio information such as track numbers. Usually, AES/EBU is sent via XLR connectors and uses a higher-voltage reference for its digital ones compared to its zeros. Some equipment is fairly forgiving about the data format and voltage level, requiring only an adapter cable to convert between the two formats. Others, including VVT, build both in. But not all equipment is so friendly — read the user's manuals before betting the farm that your audio interconnections will work without a hitch. &

Chris Meyer is a principal at Cyber-Motion, a Los Angeles-based motion graphics boutique. Meyer also serves as manager of technical research at Roland Audio Development and teaches digital audio at the Hollywood branch of the American Film Institute.

PLUG-INS FOR ELECTRIC IMAGE

NORTHERN LIGHTS 3D-F/X SERIES & ONYX COMPUTING TREE EIAS

BY RICHARD LAINHART

The Electric Image Animation System (EIAS) is the Ferrari of 3D animation programs for the Macintosh. Among its many features is support for plug-in software modules that add enhanced functionality and exotic capabilities. Like Adobe Photoshop, QuarkXPress, and Autodesk 3D Studio, EIAS's plug-in architecture is open to third-party developers. Two adventurous software companies have taken the challenge and delivered a handful of plug-ins serious 3D artists won't want to miss. Each one adds a specific function. If you happen to need that function, the plug-in can save you a huge amount of time and effort.

Northern Lights 3D-F/X Series (comprising Big Dipper 2.0, Zeus 2.0, and Dante 1.0) and Onyx Computing Tree EIAS are the only EIAS plug-ins currently shipping, although more are in the works. To use one, drag it into EIAS's SOCKETS folder and open a project. Then select it from the ADD menu and add it to the project. Generally, each plug-in's user interface consists of a single overly cluttered dialog (Dante also has one subdialog). Consequently, using them effectively often involves a lot of trial and error. This is especially true of 3D-F/X, which comes with brief disk-based manuals that don't include tutorials.

Big Dipper is the most basic of the bunch. It generates fixed or animated 3D starfields. You can set the number of stars, the size of the field and its inner radius (to create a clear area in the center), and the color temperature of the stars (based on one of seven spectral types). Things get more interesting when you tell Big Dipper to use a 3D model, say a classic five-pointed star, rather than simple points. You can also animate the starfield, rotating the entire field and each star independently — great for making them twinkle.

Zeus is ideal for generating flickering lightning and electric arc effects.

It creates 3D animated lightning bolts, providing control over color, start and end points, timing and duration, complexity (how often it breaks or changes directions), and branching (up to four forks off the main bolt). You can even add an element of randomness to the strike points. Because the bolts are 3D line segments, you can animate a camera around or through lightning strikes and apply Electric Image's deformations to create curved or warped bolts.

Dante generates particle streams to make fire, flowing water, sparks, and even pixie dust spraying from a magic wand. In addition to simple particles, Dante can use existing models, which can retain their own textures. (This is an effect you see in a lot of pharmaceutical commercials these days — textured capsules spraying from a pill bottle.) The objects in the

particle stream can change size or color over time, can respond to gravity or movement of the source object, and can be randomized within limits for more naturalistic motion.

The 3D-F/X set takes care of stars, lightning, and particle systems. Onyx's offering does trees. Tree EIAS is based on Onyx' standalone Tree Professional, which has garnered quite a reputation for its ability to generate beautiful tree models with detailed branch, twig, and leaf structures. The problem with Tree Pro models is their huge polygon count, which can reach 100,000 or more. This requires a lot of disk space to store the model and a lot of RAM to render it.

Tree EIAS dodges this problem by enabling you to open a Tree Pro tree parameters file (suffix TREE) in Electric Image, modify it using the plug-in, and render the tree in CAMERA
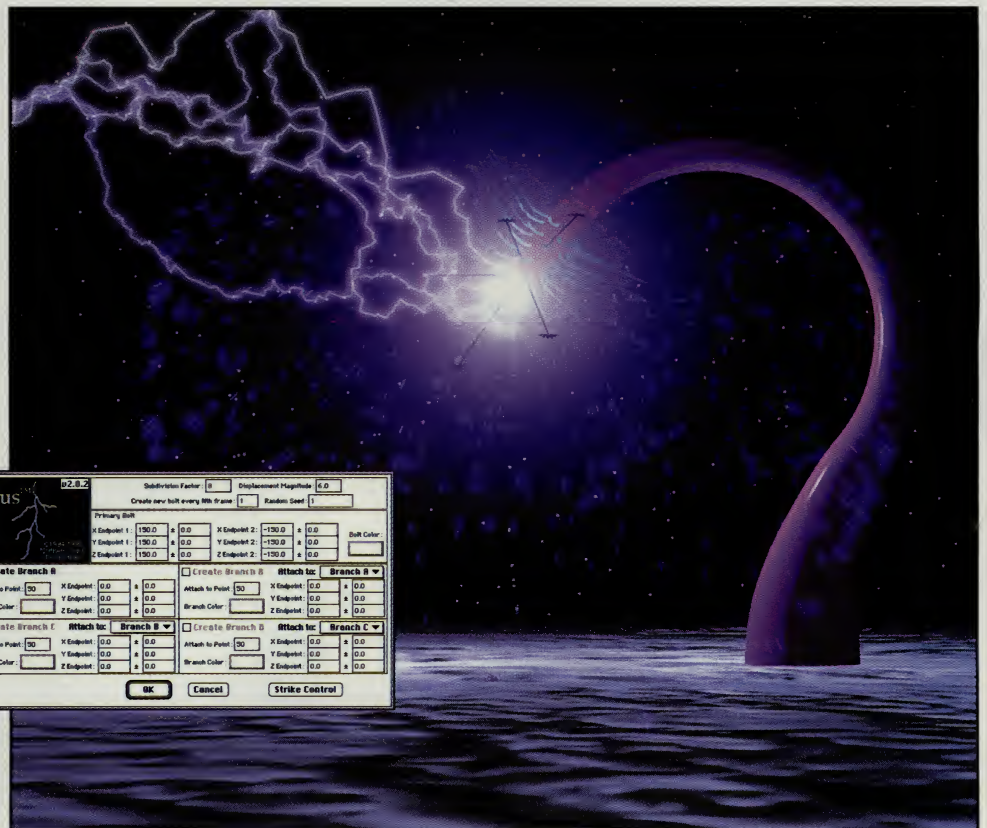
All four plug-ins in action: The background stars are from Big Dipper, the lighting is by Zeus (whose dialog screen is shown here), Dante made the sparks, and the weird tree is from Tree EIAS. The image was created entirely in ElectricImage; no external models were used.

Description

Plug-in for creating trees in Electric Image Animation System.

System Requirements

Power Mac, System 7.5 or later, Electric Image Animation System 2.5 or later, 8MB RAM beyond EIAS requirements.

Features

Imports deciduous, conifer, and palm tree parameter files created in Onyx Tree Professional; adjustable trunk, bough, branch, twig and leaf density, and image resolution; flat or curved leaves; low-res abstract of tree model reduces screen redraw time.

File Support

FACT, Image, QuickTime, PICT, TREE (Onyx Tree Pro parameters).

Suggested Retail Price

$395 ($595 with Tree Professional).

Contact

Onyx Computing, 10 Avon St., Cambridge, MA 02138; vox 617.876.3876; fax 617.868.8033; Internet onyxtree@world.std.com, http://world. std.com/~onyxtree.
Reader Service #141

Pros

Reduces memory requirements for complex tree models, makes beautiful realistic trees.

Cons

Trees begin as Onyx Tree Professional parameter files — you can't create custom trees from scratch.

Bottom Line

Realistic trees and lots of them, no sweat.

without having all those polygons onscreen. You can't create a tree from scratch as you can in Tree Pro, but you get more than 100 parameter files, including a variety of conifer, deciduous, and palm trees, plus fantasy growths like the Klingon tree illustrated here. From there, you can adjust the tree's degree of detail (leaving out structural elements such as twigs and leaf stems if the tree will be viewed only at a distance in which case they wouldn't be visible anyway) and resolution (again, for trees that won't be viewed close up). This, of course, has a corresponding impact on the final polygon count. You also have the very helpful option of displaying trees as low-res placeholders in EIAS until you're ready to render, which saves time while you're placing trees and waiting for the screen to redraw. When rendered, the realism of these trees is stunning.

Each of these plug-ins has its place. Tree EIAS is especially useful for creating realistic landscapes, while the 3D-F/X Series can create some beautiful special effects. Considering EIAS's $7,500 price tag, they're worth every penny. ☸

Richard Lainhart is a digital multimedia artist and first science officer with Fischer Multimedia Arts in New York City.

Description

Plug-ins for Electric Image Animation System that generate starfields (Big Dipper), lightning (Zeus), and particle streams (Dante).

System Requirements

Power Mac, System 7.5 or later, Electric Image Animation System 2.5 or later, 8MB RAM beyond EIAS requirements.

Features

Big Dipper 2.0: 3D starfields made up of points or user-specified models, adjustable number and color of stars, animation of individual stars, and entire starfield with random variations. **Zeus 2.0:** 3D animated lightning; adjustable color, branching, complexity, origin and strike points, start and end times, and strike speed with random variations. **Dante 1.0:** 3D particle streams made up of points or user-specified models, adjustable number of particles, start and end times, direction and speed, color, gravity, viscosity, rotation, alignment, and scaling over time with random variations.

File Support

FACT, Image, QuickTime, PICT.

Suggested Retail Price

Big Dipper, $150; Zeus 2.0, $150; Dante 1.0, $150.

Contact

Northern Lights, 210 The Village, Ste. 203, Redondo Beach, CA 90277; vox 310.376.4266; fax 310.372.5236; Internet bear60@aol.com.
Reader Service #142

Pros

Physics-based plug-ins simulate natural phenomena and create effects that would be difficult or impossible to animate by hand.

Cons

Minimal interfaces, skimpy manuals, no tutorials.

Bottom Line

Useful additions to EIAS, especially Dante.

2,000 QuickTime Clips in a Single Blow

TERRAN MOVIE CLEANER PRO 1.1.1 FOR MAC

BY RICHARD LAINHART

I t's a long-standing problem in multimedia: How do you optimize Cinepak movies for CD-ROM, especially at low data rates, and how do you do it quickly? You can Cinepak-compress captured video with almost any QuickTime-capable application, but this allows little control over quality and generally doesn't yield the best possible results. Adobe Premiere and Adobe After Effects provide a little more control, but until recently neither was able to process clips in a batch. Premiere 4.2 includes a new plug-in, CD-ROM Movie Maker, that lets you adjust compression settings and optimize Cinepak movies, but we've heard reports of problems with the current version, including unusually long rendering times. Meanwhile, After Effects makes nice-looking Cinepak movies, but it lists for $995 and doesn't provide low-level access to compression settings.

Apple's MovieShop, included with the QuickTime Developer's CD, was the first (and for years the only) program that let you tweak compression parameters and batch-compress a group of clips in a single stroke. Unfortunately, it's not a commercial product — there's no tech support, it's hard for anyone other than licensed developers to get, and it has an interface only a programmer could love. Also, it hasn't been updated in more than two years. Radius announced an updated commercial version of MovieShop called Radius Cinepak Toolkit last summer ($495 list price), but as of this writing it hasn't shipped. Unfortunately for Radius, it may be too late. The best tool yet for making Cinepak movies has just hit the market, and it costs $129.

Movie Cleaner Pro 1.1.1 from Terran Interactive is a compact, easy-to-use Power Mac utility for creating optimal QuickTime movies using any QuickTime-compliant codec. It offers powerful features for tweaking the compression process, but it's also designed to deliver excellent results to inexperienced producers. The results are spectacular. Movie Cleaner's compression routines produce the cleanest Cinepak movies I've ever seen, especially at low data rates. (A $15 shareware version, Movie Cleaner Lite 1.1, can be downloaded from Terran's Web site. It doesn't process batches and isn't native to Power Mac, so compression takes a lot more time.)
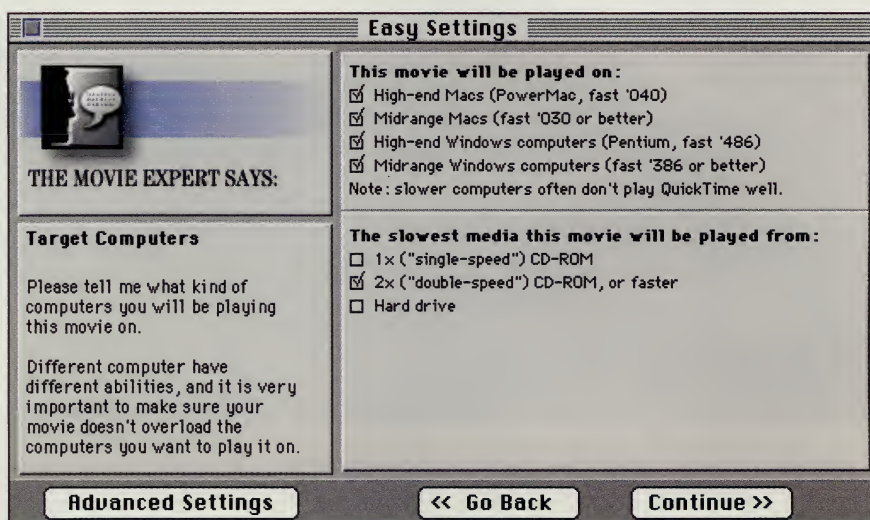
Figure 1. Movie Cleaner's Easy Settings mode leads you through a series of windows that ask questions, in nontechnical terms, about the clip's content, the way in which you intend to use it, and the delivery platform. Each set of questions is accompanied by informative text that explains why the answers matter.

The Easy Way/The Hard Way

When you launch the program, you can select either EASY SETTINGS or ADVANCED SETTINGS as your starting point. When you choose the former, Movie Cleaner prompts you to open a clip and then leads you through a series of windows that ask questions, in nontechnical terms, about the clip's content, the way in which you intend to use it, and the delivery platform (Fig. 1). Each set of questions is accompanied by informative text that explains why the answers matter. Once you've made your choices, simply click the COMPRESS button in the MOVIE CLEANER STATUS window, choose where to save
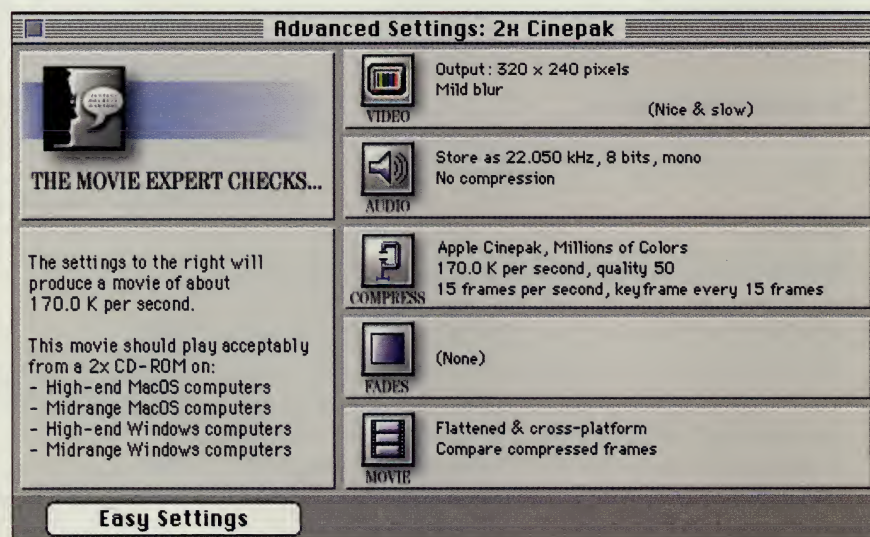
Figure 2. The Advanced Settings enable you to crop and resize; blur to reduce noise and harshness; apply fade-ins and fade-outs to video, audio, or both; attach a custom palette; and flatten for crossplatform playback.
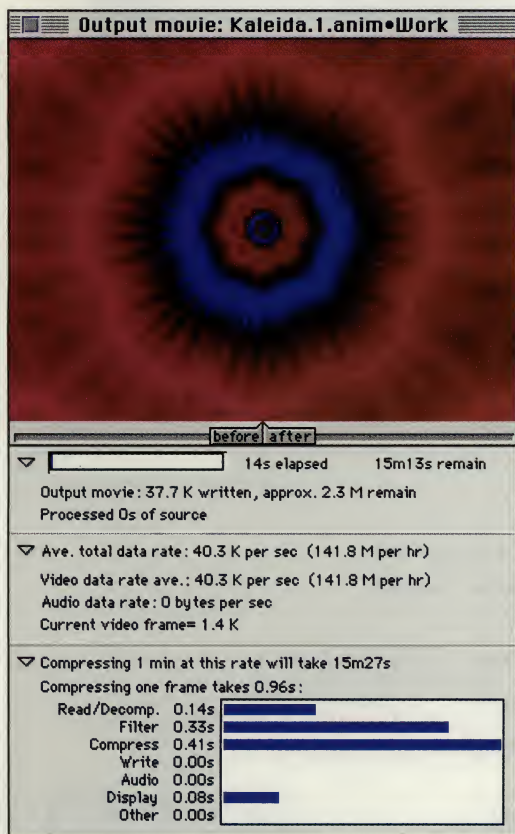
Figure 3. While compression is in progress, you get a split screen that displays the original and compressed versions of your clip side by side so you can evaluate the resulting image quality. Readouts of compression ratio, data rate, and compression time appear below.

the new movie, and let the program go to work.

While compression is in progress, you get a split screen that displays the original and compressed versions of your clip side by side so you can evaluate the resulting image quality. Readouts of compression ratio, data rate, and compression time appear below. All this is so easy that even beginners won't need to read the manual — although they should because it con-

tains a helpful overview of QuickTime and how compression works, plus a useful glossary.

Clicking the ADVANCED SETTINGS button lets more experienced users do things like crop and resize; apply a blur to reduce noise and harshness; apply fade-ins and fade-outs to video, audio, or both; attach a custom palette; and flatten a clip for crossplatform playback (Fig. 2). Beginners can learn more about the compression process by working through the easy settings and then switching to advanced mode to see what choices the program made.

The advanced settings offer various ways of cleaning up your clips and reducing their data rates, including a unique facility called the talking heads filter. Any movie that features a strong foreground element (such as an announcer), a static camera, and an unvarying background is a prime candidate. If you make a black and white (1 bit) mask that outlines the foreground element and apply the filter, Movie Cleaner will freeze the background and update only the area within the mask, maximizing the foreground image quality without increasing the data rate. Unfortunately, you must make the mask in an external program such as Photoshop, which involves copying a frame from the clip and painting over it, so it's not the most intuitive process in the world.

Batch Processing

Once you've developed useful settings, you can save them and apply them to as many as 2,000 clips at a time. Movies can be opened into a batch list one at a time (in the order you want them processed) or added en masse by dragging a folder full of clips onto the Movie Cleaner icon in the Macintosh Finder. You can process a batch using a single set of settings, or you can assign different saved settings to each clip as you enter it into the list. Since you can include the same clip in a batch list more than once, it's easy to experiment with various settings to figure out what works best for the task at hand. A cool productivity feature: You can suspend Movie Cleaner Pro's operation at any time to free your computer for other tasks, then set it to work again. It picks up right where it left off.

The interface has a couple of inconsistencies — check boxes that should be radio buttons, for instance — and Movie Cleaner is

pretty weak in the audio department. It doesn't handle sample-rate conversions very well, and it doesn't yet support the up-and-coming 16-bit IMA compression scheme. We're told these shortcomings will be addressed in the next release.

But that's mere quibbling. Movie Cleaner Pro 1.1.1 makes the best-looking low data rate Cinepak movies I've ever seen, lets you batch-compress movies with a variety of helpful tweaks, and doesn't crash in the process. Anyone producing QuickTime movies for CD-ROM needs this program. 🐾

Richard Lainhart is a digital multimedia artist and first science officer with Fischer Multimedia Arts in New York City.

Hold It Right There!

An Image Stabilization Primer

BY CHRIS MEYER

ometimes when you launch a new program, the features seem too exotic to be useful. Then, after you've worked with them for a while, you can't imagine how you lived without them. That was my experience recently with image stabilization and motion tracking.

An unsteady camera means an unsteady image, but not many organic life forms can hold a video camera perfectly steady. A number of mechanical contraptions have been invented to work around this, including the tripod, dolly, and Steadicam. No doubt these devices are a vast improvement over a pair of human hands, but even so you can end up with bobbling when you track a moving subject or zoom while trying to keep the subject centered. With a full-screen image this bobbling is often bearable, but it can cause a lot of trouble if you frame a subsection of the image (which makes any motion much more noticeable) or if you want to composite synthetic objects convincingly into real footage.

To eliminate unwanted camera motion after the fact, some people undertake the painstaking process of manually adjusting the position of each image in a clip. On most video workstations, you can accomplish this by selecting an identifiable feature of the image, such as the door handle on a car, and assigning a single anchor point to it frame by frame (or field by field). During playback, the frame itself bobbles around the anchor point, but the contents of the frame appear to stand still.

Other techniques work in special cases, such as when you're tracking objects that lack a def-inite outline, like clouds. Unfortunately, all these techniques are extremely tedious and can yield disappointing results unless you're a real pro. The adjustments required to track and stabilize an image precisely are usually on the order of fractions of a pixel.

Over the last couple of years, automated image stabilization has found its way into some of the higher end compositing programs. These systems can track an object (that door handle, for instance) and generate anchor point offsets automatically. A tantalizing example is Discreet Logic's Flame for SGI machines.

But the really good news is that automated stabilization for the rest of us has finally arrived
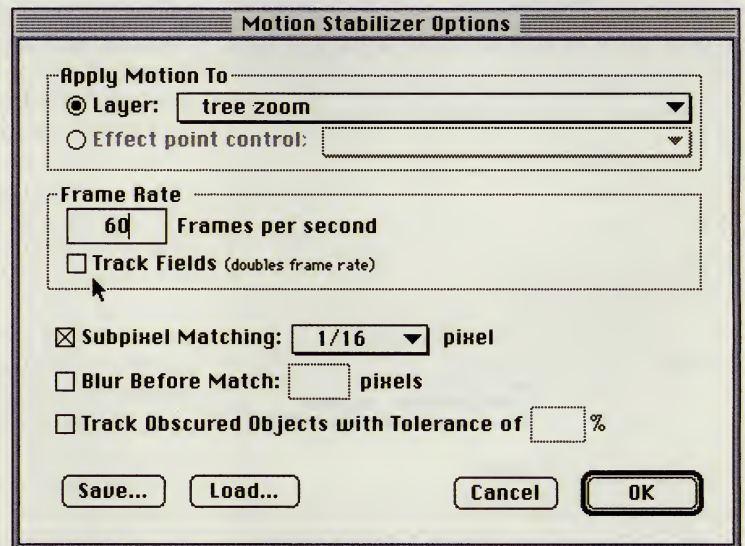
Figure 1. Setting up the stabilization options in Adobe After Effects.

Figure 2. Defining the area to track. The inner rectangle defines the portion of the image to be tracked. The outer rectangle defines the area in which the software looks for it.

Chris Meyer is manager of technical research for Roland Audio Development. A self-avowed QuickTime Baby who bought a Video-Spigot, Hi-8 camera, and Premiere 1.0 as soon as they shipped, he also serves as resident tech for his wife's desktop motion graphics company, Cyber-Motion, in Southern California.

Forget Gas, Food & Lodging

On the Information Superhighway this is the only stop you'll need.

MACWORLD EXPO/Boston
August 7-10, 1996
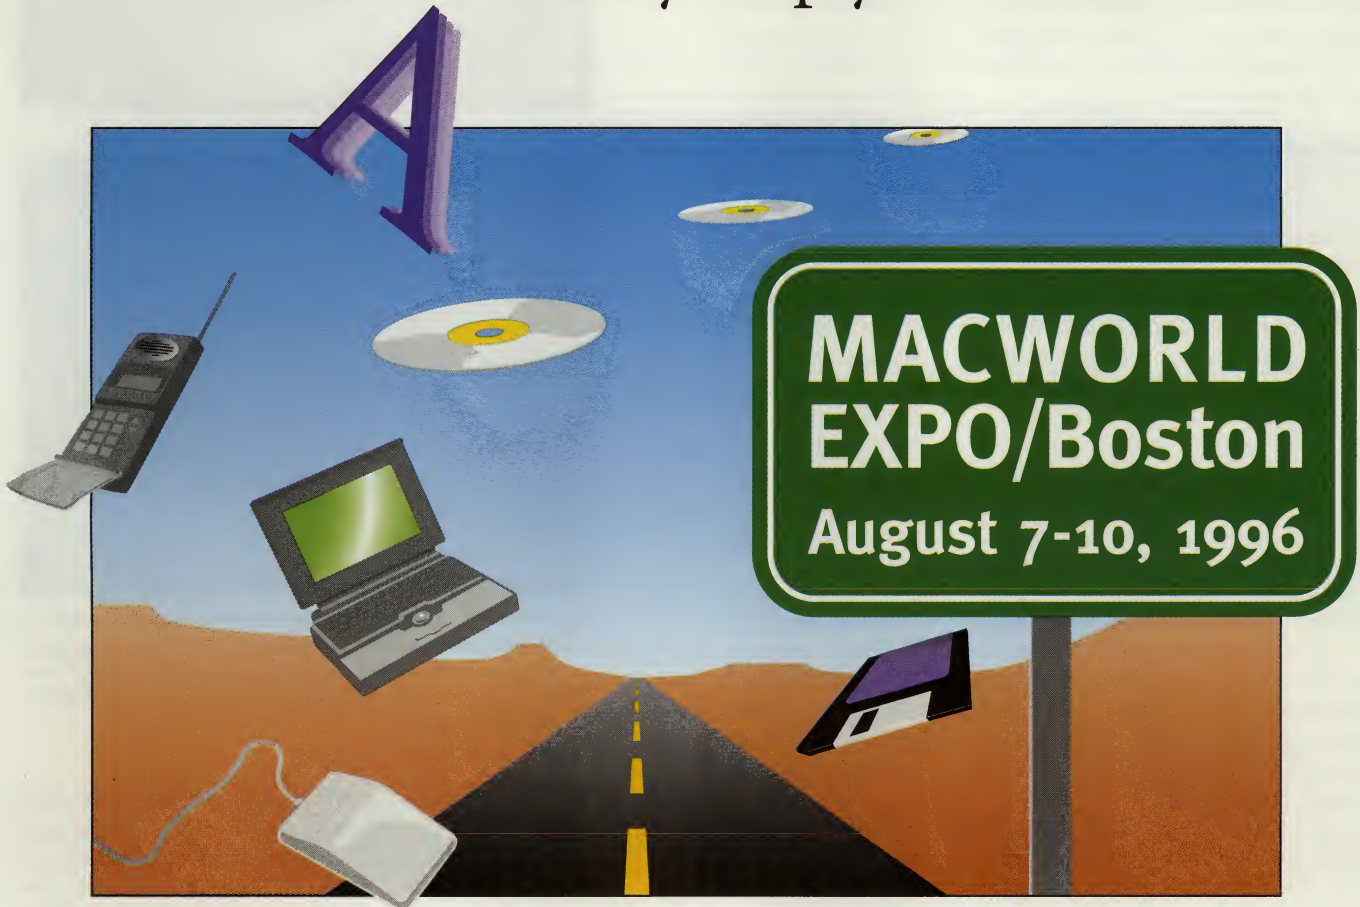
Don't want to be bypassed on the Information Superhighway? Then plan a detour to MACWORLD Expo. Here you'll test drive the products and services that enable you to maximize the potential of the Macintosh now and down the road.

MACWORLD Expo is your chance to see hundreds of companies presenting the latest in turbo-charged Macintosh technology. Make side by side comparisons of thousands of Macintosh products. Learn from the experts how to fine-tune your system and what products will keep your engine running smooth. Attend a variety of information-packed conference programs that provide the skills and knowledge to put you in the driver's seat. So pull on in and take that new Mac for a spin.

Add a stop at our upcoming **MACWORLD Expo/Boston,** August 7-10, 1996 at the World Trade Center and Bayside Expo Center to your information roadmap. We're just around the bend.

Call 800-645-EXPO
Or see us on the WWW at:
http://www.mha.com/macworldexpo

Please send me more information on MACWORLD Expo/Boston.
I am interested in: ❑ Exhibiting ❑ Attending

Name _____

Title _____

Company _____

Address _____

City/State/Zip _____

Phone _____ Fax _____

Mail to: MHA Event Management, 1400 Providence Highway,
Norwood, MA 02062   Or Fax to: 617-440-0357   ITA

FRAMES OF REFERENCE

with the release of Adobe's After Effects 3.0 Production Bundle. Along with the pro version of this Mac-based compositing program come modules for image tracking and stabilization. At CyberMotion, we've already used them to stabilize everything from a six-minute helicopter descent to literally every video fragment in a fast-paced opener for a television pilot.

What follows is a quick lesson in image stabilization for newbies (i.e., people like me just a few short months ago). I'll focus on using After Effects 3.0, but the principles can be applied to other programs as well.

Stabilization Setup and Options

Basically, there are three steps to stabilizing footage: tracking the relevant part of each frame to the anchor point, expanding and re-centering, and performing any additional processing necessary to restore the appearance of the image. With automated tracking, tracking the video to the anchor point may seem easy since the computer does the work. Actually, this step takes the most fiddling.

Let's get the setup and options out of the way first. Fig. 1 shows the STABILIZER OPTIONS dialog from After Effects. The first parameter of interest is FRAME RATE. Since high-quality captured video has interlaced fields, you'll probably want to stabilize on a field-by-field basis, especially if there's any fast or unpredictable motion. If you're working at the native frame rate of the captured footage, then check the box labeled TRACK FIELDS to have it look at individual fields as well. If you're working at the field rate rather than the frame rate (as in the example illustrated), then you don't need to check this box.

Next is SUBPIXEL MATCHING. Objects rarely move in tidy whole numbers of pixels per frame or field. Good image compositing programs work at a resolution finer than one pixel, which ensures good anti-aliasing and smooth motion. Likewise, After Effects can track objects (and determine their offset locations) with subpixel accuracy. You can adjust this resolution to various fractions of a pixel, though higher resolution makes for longer processing time. I use 1/16 of a pixel for less-than-hypercritical work. Many people like to use 1/64 (the maximum is 1/256).

The BLUR BEFORE MATCH option helps After Effects track objects. You might think sharper objects would be easier to track than blurry ones, but objects that vary between sharp and blurry can trick the matching algorithm. In a given field, objects and their edges can appear sharper than average when, through coincidences of camera tracking or motion, they align precisely with the borders of individual pixels rather than falling across them. If you have trouble tracking images, try turning this on. The image will be blurred slightly (for the

purposes of tracking only, not in the final image) to make objects more consistent from frame to frame. The cost, however, is substantial additional processing time.

Finally, some stabilization algorithms have the ability to track an object that becomes temporarily obscured, such as a moving car that becomes obscured as it passes behind a signpost. With the TRACK OBSCURED OBJECTS option enabled, if After Effects can't locate an object it's supposed to be tracking in a given field, it will try to pick it up in a later field.

Defining a Needle in a Haystack

With the options out of the way, it's time to decide what part of the frame to track. Try to identify an object or feature that's different from the area immediately around it. The more distinct it is, the less likely the software will confuse it with the background. If no zooming takes place while you're tracking, you can pick a feature almost anywhere in the frame. However, if you're trying to track an object while the camera is zooming in or out, make sure you pick an area near the center of the frame (or where you think the viewer's eye will focus). The farther away from the center of the frame, the faster the object will appear to move away from the center as the zoom progresses. If you try to stabilize by tracking an object that's even moderately off center, it may end up remaining stationary while the natural center is zoomed right off the screen.

For example, while shooting stock footage for backgrounds, we zoomed in on an interesting bark pattern on a tree. The camera was handheld, so the zoom wandered a bit — a perfect candidate for stabilization. I defined the part of the frame I wanted to track by enclosing it within the inner rectangle (see Fig. 2). There's a dark crease in the surface of the tree that seemed to hold the most promise, but since it was not in the center of the image, the zoom looked unnatural in the stabilized footage. After I increased the tracking area to include a secondary crease just below the center, the image stabilized nicely. These areas of the image may seem ill-defined, but After Effects tracked them without much trouble.

The outer rectangle tells the algorithm how widely to search to find the part of the image you've defined by the inner rectangle. If that part

Fig. 3a

Fig. 3b
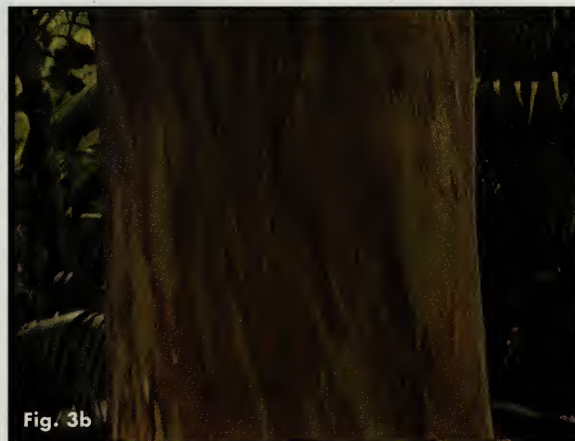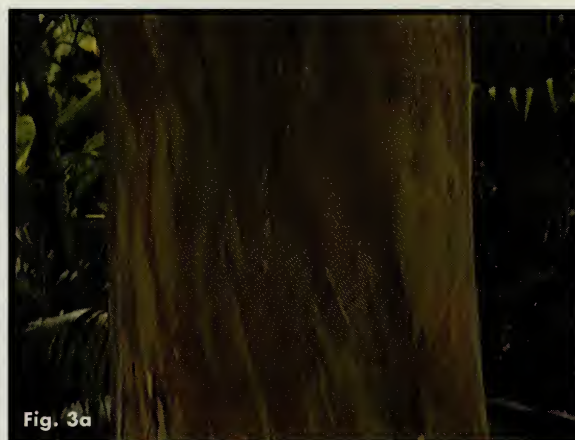
Figures 3a-b. Once a frame has been stabilized, areas of the frame may momentarily become empty as the original image jiggles around. If you plan to use only a small part of the image, this isn't a problem. Otherwise, the simplest fix is to enlarge it.

moves beyond the outer rectangle, the algorithm won't recognize it. Keep in mind that the larger the search area, the longer the search will take and the greater the likelihood that the software will "recognize" the wrong object. The outer rectangle is centered again after each frame (or field) is searched, so it only needs to be large enough to catch movement from frame to frame (or field to field), not movement over the length of the clip.

Stabilize, Expand & Center, Sharpen

After a few trials and errors — setting the tracking area, invoking the STABILIZE command, viewing the result, re-adjusting the tracking area, and stabilizing again — I had a nice, stable zoom in toward the tree. However, a new problem appeared: Since the entire image was jiggling to keep the tree steady, parts of the original image were being jiggled outside the frame. Conversely, other areas of the frame momentarily became empty. (Look for variations in the black border in Figs. 3a and 3b). If you plan to mask and use only a small part of the image, this isn't a problem. But if you intend to use the entire frame, then you'll need to do some covering up.

The simplest way to do this is to enlarge the image slightly so it always covers the frame. How much enlargement is required varies; an additional 5% to 10% usually is enough (105% or 110% of original size; type 105 or 110 into After Effects). Then re-center the expanded image to cover the frame evenly. In extreme cases, you might need to re-animate the original camera move as well. For example, in the helicopter shot I mentioned earlier, the helicopter did not just fly straight in toward the building. It also performed a 90-degree turn, wobbling, pitching, and yawing the whole way. As a result

of stabilizing the building, the turn was stabilized out of the shot as well. I added the turn back in by rotating the stabilized shot a smooth 90 degrees. When you're finished, preview each frame to make sure the image stays within the frame.

Expanding the size of an image beyond 100% reduces its resolution and may introduce a blur. To counter this effect, you might apply a filter to sharpen or unsharp-mask. The former aims to sharpen every detail in the image; the latter attempts to work on edges while leaving the rest of the image alone. Don't overdo it. Sharpen filters can accentuate video noise. Figures 4a-c illustrate the video fields originally shown in Figs. 2 and 3 after being stabilized, zoomed, and sharpened.

Note that the order in which your video application applies various processes can affect the results. For example, After Effects performs filtering before scaling, so you end up scaling a sharpened image. (Users of After Effects can cure this by performing the sharpening in a later composition.)

But Does It Do Laundry?

These capabilities can be a real lifesaver, but be aware that there are problems most image stabilization algorithms can't fix. One is the blurring imparted by extreme camera motion, which persists even after the image has been stabilized. Look at the two images in Fig. 5, which represent the first field of successive frames. Though the tree's position is stable, it appears out of focus in the second frame because the camera was jiggled sharply.

Most stabilization algorithms can compensate for rotation as well as movement along the X and Y axes. However, they often have trouble with motion and bobbling that occur in three di-
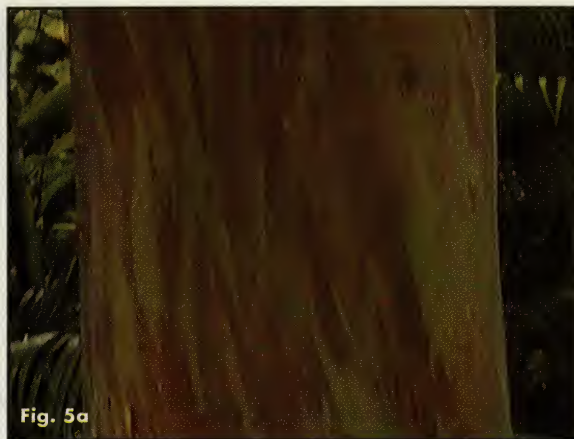
Fig. 5a
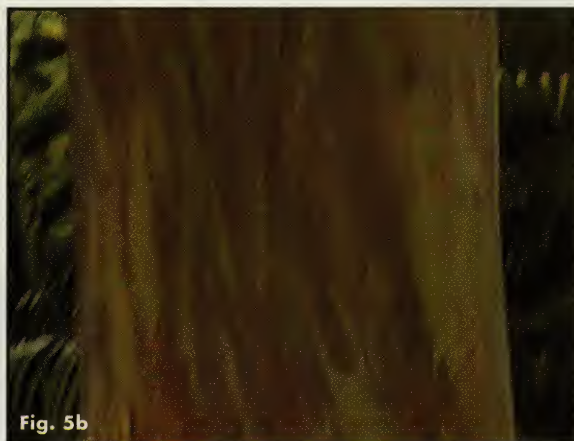
Fig. 5b

Figures 5a-b. Fast motion can introduce variable blurring from frame to frame even though the image itself has been stabilized.

mensions rather than along a line or plane. For example, a helicopter might be pitching and yawing, causing all four corners of the horizon to move up and down unequally. Stabilizing this sort of subject requires tilting the image in three dimensions, which is beyond the scope of most algorithms (at least for the present moment).

On the upside, a very useful offshoot of image stabilization is motion tracking. Rather than using anchor points to center a bobbling image, anchor points can be used to move a composited object in tandem with an object in the source footage. This technique can be used to place an animated character in the hand of an actor, replace the sign on the side of a passing bus, or composite objects into a scene so they appear to have been there when the scene was shot — bobbling and all. (By the way, adding motion blur to synthetic objects, with a shutter angle that matches that of the original camera, further enhances realism.)

On the other hand, all the wonderful creative options aside, I'm happy just to be able to use image stabilization to fix bad camera work, especially my own. I wonder how I ever lived without it. 👁
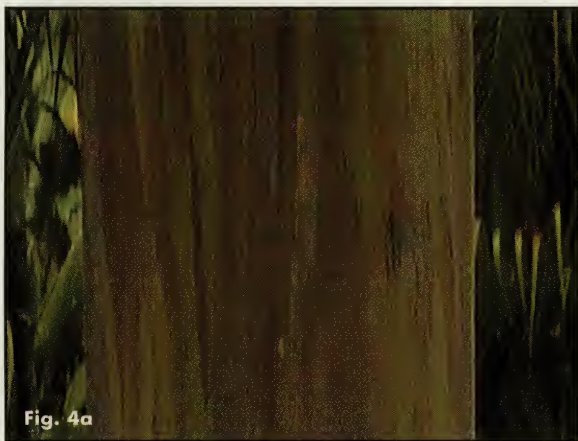
Fig. 4a
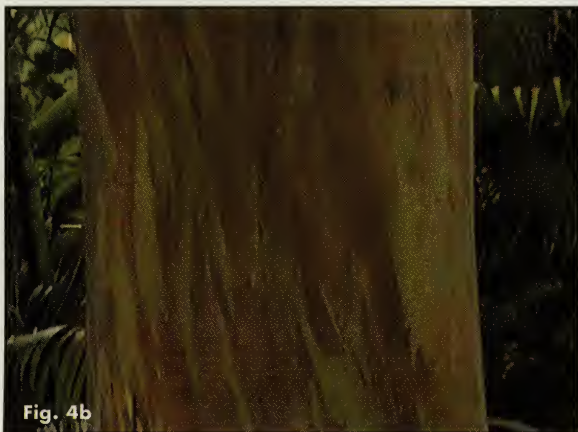
Fig. 4b
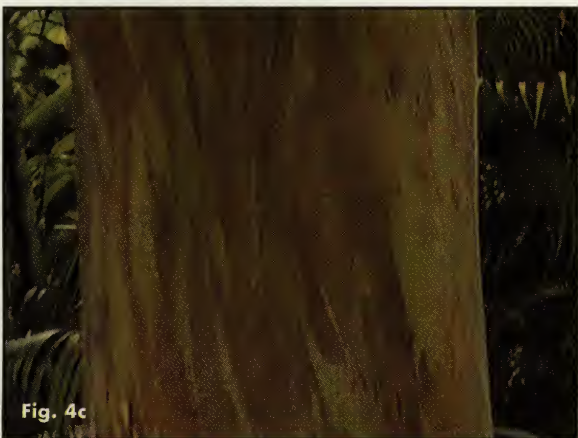
Fig. 4c

Figures 4a-c. The video fields shown in Figs. 2 and 3 after being stabilized, sharpened, and zoomed.

Mouse ≠ Cursor, Part 2

Fashioning Practical Relationships
Between the Cursor and the Mouse

BY ERIC JUSTIN GOULD

When I was a kid, I loved the story of Peter Pan. Peter showed me I could fly, Tinker Bell touched me with her magic and vulnerability, and Captain Hook left me with a lifelong curiosity about prostheses.

I'm surprised by the extent to which these characters continue to inspire me. I still believe I can fly, but have only forgotten how. I'll always have a special place in my heart for Tink, and to this day I'm fascinated by trying to imagine what it might feel like to have a metal hook for a hand.

This month's flowery illustration is of Captain Cursor (see Fig. 1). Every day, those of us who sit in front of bit-boxes don the de facto prosthesis of the desktop generation. As Captain Hook might have put it, ya see, matey, we've all got hooks fer hands.

Last month we looked at the hardware and software behind the human/machine relationship that creates this feeling. We cracked open a bunch of plastic — mice, mousepads, trackballs, trackpads, trackpoints, and joysticks. We got out our compasses and measured the landscape — global coordinates for the screen, local coordinates per window, and independent origins at the top-left corner of each.

We ascertained screen position and elapsed time to make basic calculations — distance is change in position, velocity is distance coverage over time, acceleration is change in velocity over time.

Finally, examining the use of these calculations by the operating system, it became clear that motions of the input device don't really have a one-to-one causal relationship with those of the cursor on the screen, even though they *feel* as though they are one and the same.

Eric Justin Gould is the principal of MONKEYmedia, a design and production studio in San Francisco. He is a dancer and a musician, holds degrees in cognitive science and interactive telecommunications, and is on the faculty at San Francisco State University.

"Bad Interface!"

Moving the cursor in any way other than in accordance with the intentions of the input device operator is considered one of the cardinal sins of human/computer interface design. And more often than not, the intentions of the operator are best supported by a cursor that moves with the movement of the input device.

This rule of thumb stands on solid ground. For one thing, people shouldn't be confused by

Figure 1. Captain Cursor, wielding the prostheses of the desktop generation.

things with which they choose, or are forced, to interact. Neither should they be annoyed by them. Good designs support people. All too often, computers engender both confusion and annoyance.

One of the main reasons living creatures can get through life is their amazing ability to construct a spatially and temporally continuous model of the world. If I roll a ball across the floor, I can safely assume it will move in a more or less straight line away from me. If it rolls under a

chair (and isn't obstructed by wayward cat toys), I can expect it to come out the other side. More to the point, if I put the ball in my hand, I can safely assume it will stay there and come with me as I move around the room. In any case, rarely do we expect physical objects to break out of Euclid's space-time continuum.

For this reason people have faith in the position of the cursor on the screen. The linkage between the movement of the mouse and the movement of the cursor is extremely precious. And people are sensitive to even the slightest glitches in feedback.

Fortunately, this trust and sensitivity is well placed. The cursor is given such high priority by the operating system that it takes a very serious error to sever it from the participant. When the system goes south, it isn't until the cursor freezes that we give up all hope of recovery. And when the computer isn't crashing, the cursor update almost always is given priority over other events. For the cursor's inviolate status, we owe thanks to the ancient champions of interface design.

The Changing of the Code

On the one hand, messing with the reliability of the cursor is a dangerous endeavor — if you screw up, you probably won't be forgiven. On the other hand, it can be a fruitful way to create interactions that are both useful and fun.

Sometimes the participant's intentions conflict with a direct relationship between the input device and the cursor. This can occur when the flow of either the media or the participant's motion is more crucial than the reliability of the cursor. At times like these, we're challenged to step back and question the dogma, design and build prototypes, and hire nonpartisan experts in usability testing to attempt to prove our brilliant solutions inept.

This month we'll take a code-walk through the development of a simple interaction that decouples the cursor from the mouse. The example code is written in Macromedia Director 4.0 Lingo and uses a single freeware XObject. Code listings appear later, organized by cast member.

Let's take a piece of productivity software that needs to have one of its interaction devices fixed and build a prototype of a better alternative in Director. (By showing a favorable

methodology, maybe we can persuade the publisher to implement the fix for real.) The device on the dissection table is Adobe Premiere's jog/shuttle controller for digital video and audio (see Fig. 2).
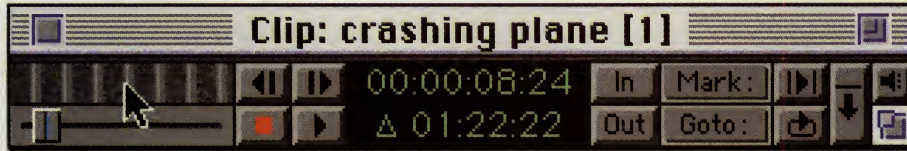

Figure 2. Adobe Premiere's jog controller in a collapsed clip window.

Jog and shuttle functions are usually paired, but they're very different. Jogging is slowly scrubbing through a clip frame by frame, while shuttling is playing a clip forward or backward at a variable speed. Many people get the two confused, so remember: Jogging is slow running, and shuttling is getting somewhere far away quickly.

To jog in Premiere, video editors press in the controller area and drag left or right to move backward or forward through a clip. As they drag beyond the bounds of the controller, the clip continues to jog until the cursor reaches

Figure 3. Jog controller feedback strip (PICT ID = 1006).

the edge of the screen. At this point, if they want to jog more, they must release the mouse button, roll the cursor back to the controller (which, across a pair of large monitors, can be a long haul), reposition the cursor over the jog controller, press, and start dragging again. If they run out of mousepad in the process, they must execute an analogous, though less obtrusive, jump-back of the mouse on the physical desktop.

Infinite Movement

Video editors who use a trackball are solely limited by the software technique employed by Adobe because trackballs are unbounded devices that can roll forever. In fact, both trackballs and mice merely track change — neither is concerned with absolute position.

Editors should be able to hook the movie at the point they grab the jog control and keep their focus on the video and audio content without interruption. The interface should be the last thing to interrupt their attention. Once they've got that virtual tape in hand, they should be able to scrub every inch of it, all the way to its ends.

The interaction we'll build is simple. When editors press on the controller, the cursor disappears. They can roll the input device left and right, jogging as far as the ends of the clip without bumping up against the screen's edge. When the mouse button is released, the cursor reappears in the same place it was before it disappeared.

We'll be moving the cursor around behind the scenes, so a small piece of external code is necessary. In MacOS, I'm aware of two choices. Jeff Fischer's moveCursor XCMD works in both HyperCard and Director. Andrew Green's Move-Mouse XOBJ works in Director and may be slightly faster. (MoveMouse can be downloaded from Electric Magic's ftp site as movemouse.hqx via ftp://ftp.netcom.com/pub/em/emagic.) We'll use MoveMouse — even though it's the cursor we're moving — and wrap it and other functions in utility handlers to keep the language understandable. So the first thing to do is set up the basic Movie, Utility, and Cast scripts.

Since our code to update the cursor is called less frequently than the operating system's code to update the cursor, occasionally the cursor may jump and flash in places where it's not supposed to be. So we'll replace the operating system cursor with a surrogate cursor sprite. In the end we'll make it invisible, too, but meanwhile

it'll be helpful in demonstrating some of the techniques we're going to be using. It's also invaluable for debugging our interaction because it limits the visible cursor to the position data we actually use.

Create a simple cursor picture that looks different enough from the system cursor that you can tell the difference and put it on the stage in channel 10 — in front of all other elements on the stage. The simpleSurrogateCursor script sets up a generic cursor that mimics the system cursor.

Grab a screen shot of a collapsed version of the clip window from Premiere. Crop to the window's size, erase the jog controller area, and place it on the stage at (-1, -1) with the INK effect set to BACKGROUND TRANSPARENT. Then set the stage size to match that of the clip window (320, 48). Open Premiere in ResEdit and copy a copy of the feedback strip (PICT ID = 1006, see Fig. 3). Place it on the stage behind the clip window at (-89, 20). It should show through the blank control area you erased.

Next, we'll bound the surrogate cursor to the area of the jog controller. The boundedCursor script is an offspring of simpleSurrogateCursor, so it inherits all methods and properties that are not overridden. The bounding is accomplished via two techniques. Up and down movement along the Y axis is constrained such that the cursor can't go above or below the control area. Left and right movement along the X axis wraps around from side to side like the spaceship in *Asteroids* — when the cursor moves beyond one edge, it jumps to the opposing edge. Both techniques can be applied creatively to a variety of situations.

Now let's get the feedback strip in the controller to move along with the cursor. The jogController script maps change in cursor position to change in position of the feedback strip. It sews the two ends of the feedback strip together, in effect making a loop that can scroll infinitely. This script also should call a movie scrubbing script, but I'll leave the media control to you.

Using the jogCursor script, which is an offspring of boundedCursor, set the jog control to engage when the editor drags in the control area. And finally, hide the surrogate cursor, which we don't need anymore, to keep the interaction clean. Now no cursor is visible during the jog interaction.

An important function the jogCursor script performs is retaining the place on the screen where the cursor last appeared so it can be restored when the video editor completes the interaction. By maintaining visible continuity of the cursor's position on the screen, we maintain the user's trust.

Premiere's shuttle technique would also benefit from a mouse/cursor retrofit. In Premiere, option-click the jog control to bring up the shuttle controller. What could be improved about the way it works? How would you do it? 🐝

INTERFACE DESIGN

Code Listing 1
Movie Scripts

```
global gCursorPosition
global gSpecialCursor

-- STARTUP
on startMovie
  openxlib "MoveMouse XOBJ"
  set gCursorPosition = MoveMouse(mnew)
  set gSpecialCursor = birth(script "jogCursor", 10)
end startMovie


-- IDLE
on idle
  mUpdate(gSpecialCursor)
  updateStage
end idle
```

Code Listing 2
Cast Script

```
on exitFrame
  go the frame
end
```

Code Listing 3
Utility Wrappers

```
-- WRAPPER FOR GETTING CURSOR POSITION
-- returns a point
on getCursorPosition
  return point (the mouseH, the mouseV)
end getCursorPosition

-- WRAPPER FOR SETTING CURSOR POSITION
-- takes a point
on setCursorPosition aPoint
  global gCursorPosition
  gCursorPosition(mSetMouseLoc, the locH of aPoint,¬
                the locV of aPoint)
end setCursorPosition

-- WRAPPER FOR GETTING SPRITE POSITION
-- takes a sprite number, returns a point
on getSpritePosition whichSprite
  return point (the locH of sprite whichSprite, the locV of sprite¬
              whichSprite)
end getSpritePosition

-- WRAPPER FOR SETTING SPRITE POSITION
-- takes a sprite number and a point
on setSpritePosition whichSprite, aPoint
  set the locH of sprite whichSprite to the locH of aPoint
  set the locV of sprite whichSprite to the locV of aPoint
end setSpritePosition
```

```
-- WRAPPER FOR GETTING X FROM A POINT
-- Lingo's H (horizontal) is confusing
on X aPoint
  return the locH of aPoint
end X


-- WRAPPER FOR GETTING Y FROM A POINT
-- Lingo's V (vertical) is confusing
on Y aPoint
  return the locV of aPoint
end Y


-- WRAPPER FOR HIDING CURSOR
-- to hide the cursor, set the cursor to 200 (blank cursor)
on hideCursor
  cursor 200  -- blank cursor
end hideCursor


-- WRAPPER FOR SHOWING CURSOR
-- to show the cursor, set the cursor to -1 (pointer cursor)
on showCursor
  cursor -1  -- pointer cursor
end showCursor


-- WRAPPER FOR HIDING A SPRITE
on hideSprite whichSprite
  set the visible of sprite(whichSprite) = FALSE
end hideSprite


-- WRAPPER FOR SHOWING A SPRITE
on showSprite whichSprite
  set the visible of sprite(whichSprite) = TRUE
end showSprite
```

Code Listing 4
simpleSurrogateCursor

```
property pSprite  -- which sprite

on birth me, whichSprite
  hideCursor
  set pSprite = whichSprite
  puppetSprite pSprite, TRUE
  showSprite(the pSprite of me)
  updateStage
  return me
end birth


on mUpdate me
  set cursorPos = getCursorPosition()
  set cursorPos = mModify(me, cursorPos)  -- special rules
  setSpritePosition(pSprite, cursorPos)
end mUpdate


on mModify me, cursorPos
  -- no special functionality by default
  -- offspring can override this method with unique cursor rules
  return cursorPos
end mModify
```

Free Info! Fast!

Free Product Information from *InterActivity* magazine.

Get valuable facts about products and services that appear in *InterActivity*.

Return the Card Today!

There's no cost and never any obligation to you!

Here's all you do: Write your name and address on the attached card, complete the brief survey, and then circle the numbers on the card that match to ones at the bottom of ads or are shown on the following page. Then fax or mail the card to us TODAY. We even pay the postage!

FAX NUMBER
413-637-4343

InterActivity
THE HOW-TO MULTIMEDIA MAGAZINE

Free Product Information from *InterActivity!* (Please Print)

For faster service, fax this card to 413-637-4343

A. When do you plan to purchase?
- ☐ 1. Now
- ☐ 2. Six months
- ☐ 3. 12 months
- ☐ 4. Only gathering information

B. On which platform do you develop?
- ☐ 5. Windows
- ☐ 6. Macintosh
- ☐ 7. DOS
- ☐ 8. UNIX Workstation
- ☐ 9. OS/2
- ☐ 10. Other_____

Name _____
Title _____
Company _____
Address _____
City _____ State/Province _____
Country _____ Zip/Postal Code _____
Phone _____ Fax _____
Internet/Email Address _____
May we contact you via Email?   ☐ Yes   ☐ No

ISSUE DATE: APRIL 1996					EXPIRATION DATE: JULY 31,1996								
1	19	37	55	73	91	109	127	145	163	181	199	217	235
2	20	38	56	74	92	110	128	146	164	182	200	218	236
3	21	39	57	75	93	111	129	147	165	183	201	219	237
4	22	40	58	76	94	112	130	148	166	184	202	220	238
5	23	41	59	77	95	113	131	149	167	185	203	221	239
6	24	42	60	78	96	114	132	150	168	186	204	222	240
7	25	43	61	79	97	115	133	151	169	187	205	223	241
8	26	44	62	80	98	116	134	152	170	188	206	224	242
9	27	45	63	81	99	117	135	153	171	189	207	225	243
10	28	46	64	82	100	118	136	154	172	190	208	226	244
11	29	47	65	83	101	119	137	155	173	191	209	227	245
12	30	48	66	84	102	120	138	156	174	192	210	228	246
13	31	49	67	85	103	121	139	157	175	193	211	229	247
14	32	50	68	86	104	122	140	158	176	194	212	230	248
15	33	51	69	87	105	123	141	159	177	195	213	231	249
16	34	52	70	88	106	124	142	160	178	196	214	232	250
17	35	53	71	89	107	125	143	161	179	197	215	233	
18	36	54	72	90	108	126	144	162	180	198	216	234	

Free Info! Fast!

Free Product Information from *InterActivity* magazine.

Get valuable facts about products and services that appear in *InterActivity*.

Return the Card Today!

There's no cost and never any obligation to you!

Here's all you do: Write your name and address on the attached card, complete the brief survey, and then circle the numbers on the card that match to ones at the bottom of ads or are shown on the following page. Then fax or mail the card to us TODAY. We even pay the postage!

FAX NUMBER
413-637-4343

InterActivity
THE HOW-TO MULTIMEDIA MAGAZINE

interact with InterActivity™

COMPANY	RS#	PAGE	PHONE #
AUDIO			
InVision	110	14	415.812.7380
Opcode	112	16	415.856.3333
YAMAHA	**10**	**21**	**714.522.9011**
AUTHORING/PRESENTATION SOFTWARE			
DIGITAL DELIVERY	**8**	**15**	**617.275.3830**
Electric Magic	151	79	415.759.4100
Fischer Group	150	79	909.597.1123
INTL INTERACTIVE MEDIA	**11**	**13**	**617.290.5964**
Nichimen Graphics	111	14	310.577.0500
ORACLE	**1**	**C2**	**415.506.6577**
VIDI	114	16	818.358.3936
CD-ROM PLAYERS/RECORDERS			
COREL	**22**	**C4**	**800.772.6735**
			353.1.706.3912
CONFERENCES/ASSOCIATIONS			
3D DESIGN CONFERENCE	**9**	**19**	**415.905.4994**
COMPUTER GAME DEVELOPERS CONFERENCE	**17**	**17**	**415.905.2702**
MACWORLD BOSTON	**18**	**75**	**800.645.EXPO**
NATIONAL ASSOCIATION OF BROADCASTERS	**19**	**33**	**202.429.4188**
WEB DESIGN & DEVELOPMENT '96	**16**	**61**	**800.441.8826**
GRAPHICS/3D MODELING/ RENDERING/ANIMATION			
3NAME3D	**39**	**93**	**310.314.2171**
Ascension	113	16	802.860.6440
auto·des·sys	155	96	614.488.9777
BIOVISION	**40**	**93**	**415.292.0333**
ELECTRIC IMAGE	**21**	**C3**	**818.522.1622**
ElectricImage	156	96	818.577.1627
Macromedia	157	96	415.252.2000
Northern Lights	142	71	310.376.4266
Onyx Computing	141	71	617.876.3876
STRATA	**2**	**3**	**801.628.5210**
TAMRON INDUSTRIES	**12**	**38**	**516.484.8880**
The Valis Group	158	96	415.435.5404
MONITORS/DISPLAYS			
BREAKAWAY	**32**	**91**	**508.562.7666**

COMPANY	RS#	PAGE	PHONE #
ONLINE/INTERNET/TELECOMMUNICATIONS			
InterVista	120	20	415.434.8765
ParaGraph	121	20	408.364.7700
WebMaster	122	20	408.345.1800
RESOURCES			
CARRONADE GROUP	**14**	**39**	**213.935.7600**
INTERACTIVITY BACK ISSUES		**69**	**800.444.4881**
KEYBOARD CENTRAL	**10**	**21**	**www.keyboardmag. com/keyboard**
LIGHTWAVES	**33**	**92**	**617.523.0625**
SERVICES			
Cantwell and Paxton Inc.	154	90	703.415.1616
DISC MAKERS AVL	**38**	**93**	**609.663.9030**
MEGALODON MULTIMEDIA	**34**	**92**	**813.328.2283**
TACHYON TECHNOLOGY	**36**	**92**	**206.622.7805**
Thomson & Thomson	152	90	617.479.1600
Trademark Research Corp	153	90	212.228.4084
VANCOUVER FILM SCHOOL	**31**	**91**	**604.685.5808**
STOCK MEDIA			
COREL	**5**	**9**	**800.772.6735**
			353.1.706.3912
COREL	**6**	**11**	**800.772.6735**
			353.1.706.3912
Graphic Detail	118	20	502.363.2986
PICTURE NETWORK INTL	**4**	**7**	**703.807.2789**
THE MUSIC BAKERY	**35**	**92**	**800.229.0313**
STORAGE			
FWB	115	18	415.325.4392
MicroNet	119	20	714.453.6100
UTILITIES			
S&S	116	18	617.273.7400
VIDEO			
DATA TRANSLATIONS	**3**	**4-5**	**508.460.1600**
Play	117	18	916.851.0800
Radius	140	67	408.541.6100
TERRAN INTERACTIVE	**30**	**91**	**408.353.8859**
Terran Interactive	143	73	408.353.8859
U LEAD SYSTEMS	**20**	**87**	**310.523.9396**

INTERFACE DESIGN

Code Listing 5
boundedCursor

```
property ancestor -- to establish inheritance
property pBounds -- left, top, right, bottom

on birth me, whichSprite
  set ancestor = birth(script "simpleSurrogateCursor",¬
                          whichSprite)
  set pBounds = rect(0, 21, 78, 34)
  return me
end birth me

on mModify me, cursorPos
  if not(inside(cursorPos, the pBounds of me)) then
    set cursorPos = mBoundPointInRect(me, cursorPos,¬
                            the pBounds of me)
    setCursorPosition(cursorPos)
  end if
  return cursorPos
end mModify

on mBoundPointInRect me, aPoint, aRect
  -- wrap left and right
  if x(aPoint) < the left of aRect then
    set aPoint = point(the right of aRect, y(aPoint))
  else if x(aPoint) > the right of aRect then
    set aPoint = point(the left of aRect, y(aPoint))
  end if
  -- stop at top and bottom
  if y(aPoint) < the top of aRect then
    set aPoint = point(x(aPoint), the top of aRect)
  else if y(aPoint) > the bottom of aRect then
    set aPoint = point(x(aPoint), the bottom of aRect)
  end if
  return aPoint
end mBoundPointInRect
```

Code Listing 6
jogCursor

```
property ancestor      -- to establish inheritance
property pActive       -- whether jogging
property pStartPoint   -- place user last saw cursor
global gJog

on birth me, whichSprite
  set ancestor = birth(script "boundedCursor",¬
                          whichSprite)
  showCursor -- since simpleSurrogateCursor hid it
  hideSprite(the pSprite of me)
  set gJog = birth(script "jogController", 2)
```

```
  return me
end birth me

on mUpdate me
  set cursorPos = getCursorPosition()
  if the mouseDown then
    if (not pActive) and (inside(cursorPos, the pBounds¬
                                of me)) then
      set pActive = TRUE
      hideCursor
      set pStartPoint = cursorPos
      set the pLastCursorPositionX of gJog = x(cursorPos)
    else if pActive then
      mUpdate(ancestor)
      mUpdate(gJog)
    end if
  else if pActive then
    set pActive = FALSE
    setCursorPosition(pStartPoint)
    showCursor
  end if
end mUpdate
```

Code Listing 7
jogController

```
property pMinX, pMaxX
property pLastPositionX
property pLastCursorPositionX
property pJogSprite

on birth me, whichSprite
  set pMinX = -49
  set pMaxX = 127
  set pLastPositionX = ( pMinX + pMaxX ) / 2
  set pJogSprite = whichSprite
  puppetSprite pJogSprite, TRUE
  return me
end birth me

on mUpdate me
  set cursorPositionX = x(getCursorPosition())
  set distance = cursorPositionX - pLastCursorPositionX
  set pLastCursorPositionX = cursorPositionX
  set jogPosition = getSpritePosition(pJogSprite)
  if (x(jogPosition) + distance) > pMaxX then
    set newJogPositionX = pMinX
  else if (x(jogPosition) + distance) < pMinX then
    set newJogPositionX = pMaxX
  else
    set newJogPositionX = x(jogPosition) + distance
  end if
  setSpritePosition(pJogSprite, point(newJogPositionX,¬
                                y(jogPosition)))
end mUpdate
```

InterActivity™ Sales Team
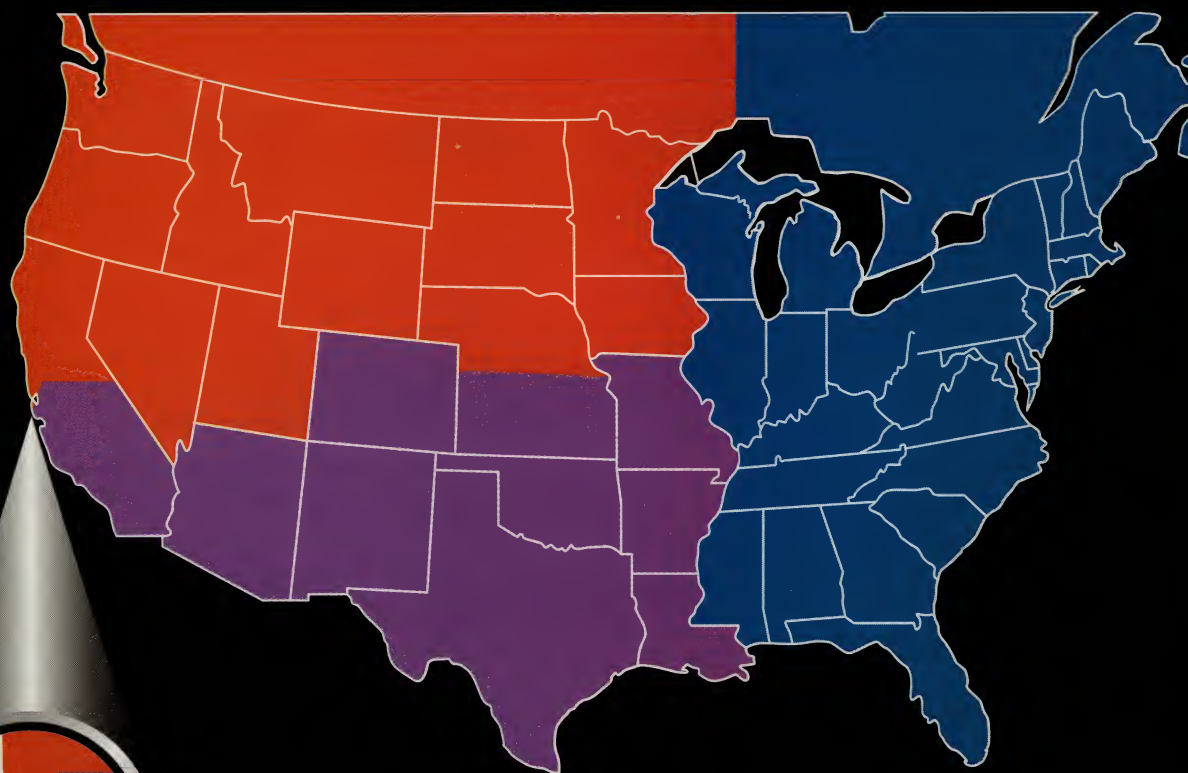
415 Area Code: Northwest Territory

408 & 510 Area Code: Southwest Territory

David Morrison, Eastern Advertising Manager
voice: 212.615.2325 • fax: 212.279.3962
email: dmorrison@mfi.com

Lisa Rahmanan, Marketplace Sales
voice: 212.615.2889 • fax: 212.279.3962
email: lrahmanan@mfi.com

Leigh Godfrey, Ad Traffic Coordinator
voice: 415.905.2422 • fax: 415.905.2236
email: lgodfrey@mfi.com

Alicia Langer, Northwest Advertising Manager
voice: 415.655.4283 • fax: 415.655.4360
email: alanger@mfi.com

James Pothier, Southwest Advertising Manager
voice: 415.655.4239 • fax: 415.655.4360
email: jpothier@mfi.com

Eleni Koros, Marketplace Sales
voice: 415.655.4281 • fax: 415.655.4360
email: ekoros@mfi.com

The Whole Enchilada

New Paradigms for Audio Production in Nonlinear Media

BY DAVID JAVELOSA

Having been ground through the mill of the typical multimedia production cycle a number of times over the past few years, I've come to approach the process differently than I did when I was composing for other media. Obviously, working in nonlinear media is very different from working on more traditional kinds of projects, but one of the most important adaptations I've made is not technological. It's in my relationships with others collaborating on a given project. Looking at these relationships differently has streamlined my work and saved me a lot of time (which, as we all know, is money) while pounding out a project under deadline and under budget.

This is a different direction than this column usually takes, focusing on process and politics rather than technology. Nonetheless, it does seem analogous to my usual territory, the complex job of balancing hardware, software, creativity, ingenuity, time, and money to deliver the audio portion of an interactive production.

A Team of One

"Doing sound" for a software production, be it a multimedia encyclopedia, a computer game, or whatever, can mean various things depending on who you're talking to. At ground level, among the actual development team, it often means

David Javelosa is a faculty member at the UCLA Extension and with the San Francisco State University Multimedia Studies Program. Formerly senior music designer for Sega Studios, his recent work includes music for Marvel CD-ROM Comics, Disney's *Pocahontas* video game, and an LP of 31st-Century lounge music.
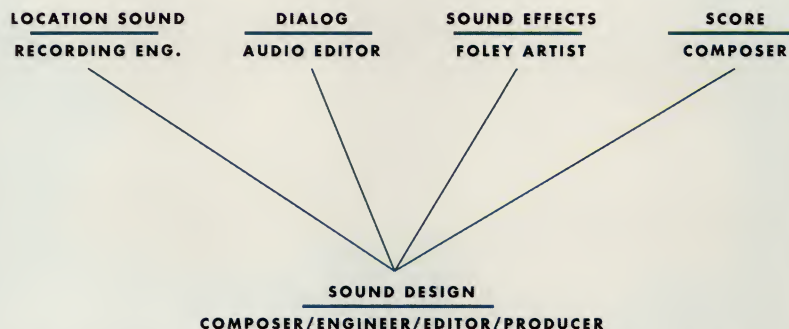
Figure 1. The role of the sound designer in multimedia. Formerly delegated to specialists, the responsibilities for location sound, dialog and voiceover recording, sound effects, and music scoring all fall on the shoulders of the multimedia audio person.

the whole enchilada. In their view software production is all-inclusive. At the very least the job includes music, instrument sounds, and sound effects. If the project is intended for CD-ROM or contains video or voiceover elements, the job includes location sound recording and dialog recording and editing. To the typical musician this can be a daunting task (as well as a cause of distress to the uninitiated producer or lead programmer).

These days, large productions have become the norm, and the garage-band development team is no longer solely responsible for securing diverse music and audio resources. Producers who have been around the block or those who have migrated from film, TV, or other linear media are familiar with the various specialists of the audio world: the composer who writes, orchestrates, and possibly conducts the score; the foley artist who records, constructs, or otherwise comes up with the sound effects; the location sound engineer who captures live sound as the video is being shot; the audio editor who makes sure the dialog sounds consistent and syncs properly with the video; and the mixdown engineer who keeps all these elements balanced.

Sounds like a big Hollywood production, doesn't it? You bet it does. And it would be a sensible approach if software producers were accustomed to burning money the way film producers do. Unfortunately, for the relatively small multimedia house or even a medium-sized game company, this model is the road to bankruptcy. In a medium-sized audio group at Sega, we had a studio engineer, a sound ef-

fects person, and a music programmer — all of whom were able to write music. This enabled us to cover the many tasks involved in producing game sound and provided some variety in the music we were able to create.

Still, in many situations the ground-level model makes a lot of sense: Assign a person to handle all audio production. This works best if the audio person knows when to contract out the work they are least able to handle well. In any case, in the world of interactive multimedia, the job of composer has mutated into something increasingly known as *sound design* (see Fig. 1).

Production vs. Postproduction

In the film and video biz, the composer's role is to "hurry up and wait." Actually, a more precise phrase is "hurry up and wait, then really hurry up or die." In a linear medium, the producer is concerned with finishing the elements in order, and music comes pretty close to last — when the accumulation of missed deadlines reaches the meltdown point. Countless meetings may have been held early on to determine what the music should sound like, and placeholder music may have been produced for rough cuts, but the producer isn't going to want to hear a single note until well into the editing and assembly phase.

By and large, nonlinear media aren't produced this way. Everything is happening at once. Scenes and characters are thrown together for demos at the drop of a hat and brought to brainstorming meetings for instant feedback from other members of the team. Music, sound effects, and even underscore for video clips may be required as

early as the design phase. Furthermore, as we've already noted, the audio person often wears many hats. If the final music isn't needed before the end of the process, other audio elements certainly are.

Properly conceived, sound design starts at the beginning and continues throughout the production process. Ironically, just as the video game technician of old controlled the production of audio, so should the sound designer coordinate all the audio elements in a multimedia production. Consistency among sound effects, score, dialog, transitional music, and so on becomes critical in an interactive environment, where the user's actions determine the flow of the presentation. In fact, the more complex the interactive experience, the greater the impact of consistent sound design.

Close Encounters

Not only do multimedia sound designers wear many hats, they must integrate their work with other members of the development team. The word "relationship" has always reminded me more of astrophysics than interpersonal romance, so let's consider the ways in which various members of a production team orbit one another.

In the traditional production model, all departments, including sound, report to the producer, who coordinates and mediates among artistic considerations. The composer deals with this person exclusively (see Fig. 2). Creative synergy is strictly a vertical phenomenon, with cross-pollination of ideas and inspiration occurring only at the top level. Although it's seldom ideal, this scheme does support a quick turnaround time for each of the elements — a dictatorship gets things done faster than a democracy. The risk, of course, is that the various specialists will view the project with tunnel vision (at best) or work disjointedly or even at cross-purposes (at worst).

In an interactive production, however, the specialists not only must be aware of one another's work and have a unified thrust, but they must also design their own elements to integrate properly within the interactive framework. An organizational structure that supports this degree of collaboration is illustrated in Fig. 3. For the sound designer, this presents a rare opportunity to tailor each element to the production as a whole. Let's examine the possible encounters the sound designer might have with other members of the interactive development team and the questions that might affect the audio person's work along the way.

▶ Composer -> Designer: For a given project, what is the definition of interactivity? Should the world we're creating have a particular sound? Do the user's actions trigger abrupt transitions, or are they gradual? Should the user's actions be responsible for changes in the sound and/or music? Is different music required each time the user repeats a particular action?
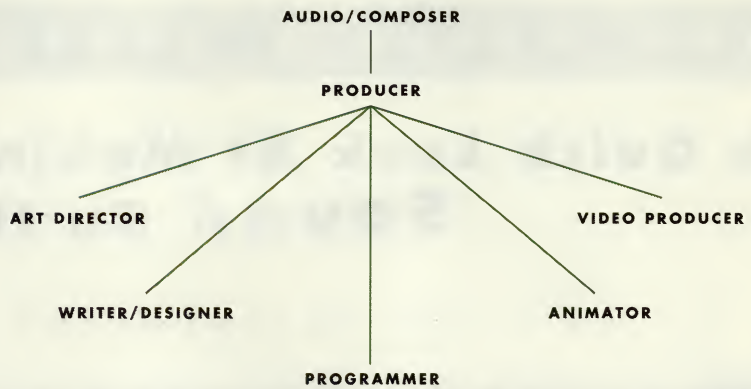
Figure 2. The traditional production model for linear media. All departments report to the producer. Cross-pollination of ideas and inspiration occurs only at the top level. This arrangement supports a quick turnaround time, but it risks lack of coordination among various specialists.

▶ Composer -> Programmer: Are special tools needed to develop sounds for this particular product or to play them on the delivery platform? How will various audio resources be used? In what file formats and on what media should they be delivered? What are the memory restrictions? Given the technical limitations of the product specification, is the sound design concept feasible?

▶ Composer -> Art Director: How should the sound design complement the product's look? Is there agreement or conflict between the visuals and the background score? Do the characters or background art suggest special sound effects?

▶ Composer -> Animator: What kinds of sounds do the characters make? What pieces of equipment do they use and what should they sound like? Should each character have a musical theme? Does anything happen that requires sound effects? Is there any rendered action that requires a synced score?

▶ Composer -> Video Producer: Will the soundtrack include location recordings, voiceover, or dialog? Are additional sound effects needed? What degree of synchronization is required? Will alternate foreign language versions be made? When the raw video is converted to QuickTime or other digital video formats, will the soundtrack need to be converted and dithered separately?

▶ Composer -> Producer: Since the producer traditionally has the final word, how will quality control be delegated in each of the preceding encounters? At what level of detail will the producer be involved in audio production? How much interaction should there be between the sound designer and other specialists? Is the sound design appropriate to the production as a whole?

The Final Word

. . . should always be the producer's. Nonetheless, interaction between the audio person and the rest of the development team can save a lot of time and effort. It can even make the difference between a mediocre product and a stellar performance — especially when the sound department is a team of one. The challenge is to make use of everyone's inspiration and ingenuity without diffusing the producer's vision. 🐾
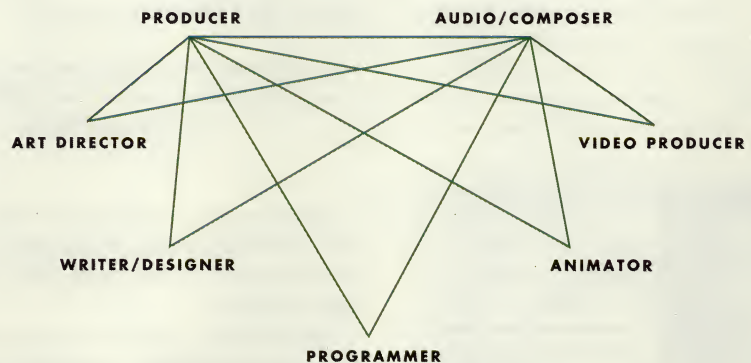
Figure 3. A production model for interactive multimedia. Specialists are in direct communication with one another, enabling them to work in a unified, tightly integrated manner. This presents a rare opportunity for the sound designer to tailor each audio element to the production as a whole.

A Quick Look At Making
Sound On the BeBox

BY JOHN WORTHINGTON

hen you think about it, it's so obvious. It's strange I even have to say it. Sound is one of the most important parts of any multimedia project. Sound adds another dimension to a presentation. It stirs an emotional reaction. If I had to choose, I'd take a slide show with sound over video or animation without sound. Even silent movies weren't so silent. There was always a musician in the theater accompanying the movie.

Oddly, sound is often an afterthought. It's hard to do well. You usually want sound to be playing in the background, and if you don't have a multitasking OS, you have to do a lot of fancy work with interrupts. Having written this sort of software twice for the Macintosh Sound Manager and Macintosh MIDI Manager, I can say it's not for the faint of heart.

BeBox the BeatBox

The BeBox has a multitasking operating system and a media kit that makes it easy to play music in the background, that is, asynchronously. At the heart of this capability is the audio server and a class called BAudioSubscriber.

All computers handle digital audio basically the same way. The digital audio, which is just a buffer of data, gets written to a digital audio converter (DAC). The DAC converts the data back to an analog signal that goes out to the speaker.

The Audio Server passes a list of buffers between the DAC and BAudioSubscriber objects. These buffers flow along a path called a stream. Think of it this way. A stream is a list of BAudioSubscriber objects that each get called in turn with a buffer. This feature makes it easy to do things like realtime effects. For example, you

John Worthington is the founder and president of MojoSoft, a company specializing in multimedia technology and development. In his misguided youth, he wrote the Macintosh MIDI Manager and was the project leader for QuickTime. His most recent project is the MusicNet CD-ROM.

could write a BAudioSubscriber object that implements reverb and insert it in the stream.

To get a BAudioSubscriber object in on the action, first it has to subscribe to the stream. Then you need to write two routines. The first is the stream function, which gets called by the audio server whenever a buffer needs processing. The second routine is the completion function. This function gets called when your BAudioSubscriber exits the stream so you can do whatever cleanup is necessary.

Let's look at some code. Here's a fragment that shows how to subscribe to an audio stream. To save space, I've omitted the error checking code. Of course, you shouldn't do this in any real code. Always check for errors. It'll save a lot of debugging time later. In this case, any time we assign
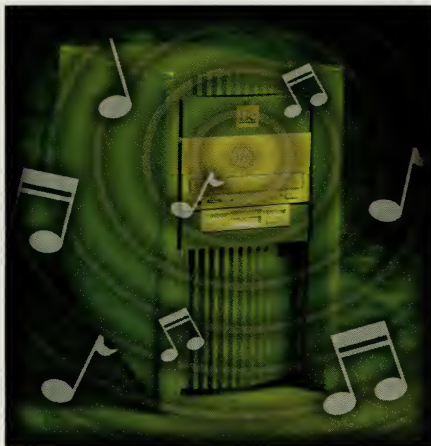
a value to a result we should check to see whether it's less than 0. If so, there's a problem.

```
mySub = new BAudioSubscriber("my sub");
result = mySub->Subscribe( DAC_STREAM,
                NO_SUBSCRIBER_ID,
                FALSE);
```

These first two lines create the object and subscribe to the stream. The string "my sub" is the optional print name for the object. It's useful for debugging.

The second line tells the new object to subscribe to the DAC's stream. Remember, you can have multiple BAudioSubscribers in a stream. For the second parameter, you either pass the constant NO_SUBSCRIBER_ID, which means you want a new stream, or the ID of an existing stream. If

the final parameter is TRUE, the function won't return until the BAudioSubscriber is successfully subscribed. Be warned though, there's no timeout. You could be waiting a long time if some other process is holding onto the DAC's stream.

The next three lines set up some of the characteristics of the stream.

```
result = mySub->SetStreamBuffers(
            kAudioBufferSize,8);
result = mySub->SetDACParameters(
            kBytesPerSample,
            kSamplesPerFrame,
            SOUND_FORMAT_BIG_ENDIAN,
            SOUND_FORMAT_LINEAR);
result = mySub->SetCODECParameters(
            44100);
```

The first line sets the size and number of buffers used by the stream. On one hand, having large buffers is a good idea because you spend less time refilling the buffers. The downside is that a longer delay occurs when stopping the sound, and when you do have to fill a buffer, it takes longer.

A brief note about the constants like kAudioBufferSize. These are constants I set up in my applications. These constants let me try different buffer sizes without worrying about changing all the places that might rely on the buffer size.

SetDACParameters tells the DAC what format the audio is in. The DAC in the BeBox supports a variety of formats. The number of bytes per sample is pretty easy. If you have 8-bit audio, it's one byte per sample. If you have 16-bit audio, it's two bytes. For monophonic sound, kSamplesPerFrame would be one. For stereo, it would be two. SOUND_FORMAT_BIG_ENDIAN says that the bytes are in Big Endian format. This really only applies when you have more than one byte per sample. With the Big Endian format, the most significant byte occurs at the lowest memory location. With Little Endian, it's the opposite. The sample format is even easier. Unless you're working with telephone audio, you'll want to set it to SOUND_FORMAT_LINEAR.

Here's the last line:

```
result = mySub->EnterStream(NULL,
            TRUE,
            this,
```

&streamFunction,
&compRoutine,
TRUE);

You'll recall that the stream can be thought of as a list of BAudioSubscribers that the buffer is passed between. Good thing you were paying attention. The first parameter is the BAudioSubscriber that we want to be next to in the stream. The order in which the BAudioSubscribers get passed buffers is very important if you're trying to do realtime effects. For example, you probably want the BAudioSubscriber that reads audio from the disk at the beginning of the stream. To place a BAudioSubscriber at the beginning of the stream, we use NULL for the first parameter and TRUE for the second. If we had wanted to be next to a particular BAudioSubscriber, we could have passed its ID here instead of NULL.

When the second parameter is TRUE, our BAudioSubscriber will be placed before the BAudioSubscriber specified or, if the first parameter is NULL, at the beginning of the stream. If this parameter is FALSE, then the opposite happens.

The third parameter is a userData field. You can put whatever you want here, and it'll be passed as an argument to the stream and completion routines. It's very important that the stream and completion routines aren't member functions of your objects. For this reason, I always pass the C++ keyword "this" as my userData. This method lets me access the object from within these routines. It's a trick worth remembering.

The fourth and fifth parameters are the addresses of our stream and completion functions. I'll get to these in a minute.

The last parameter tells the Audio Server to run the code in the background. If this is TRUE, then a separate thread is spawned for your stream function.

Stream and Completion Functions

The stream function gets called with each buffer that passes through the stream. You make whatever changes you want to the data and return. When you return, the buffer is passed to the next BAudioSubscriber.

Here's what it looks like:

```
bool streamFunction(void * userData,
        char * buffer, long count)
{
BAudioSubscriber *  subscriber;
bool    result;

    // cast user data into our object
subscriber =
        (BAudioSubscriber *)userData;

    result = RdFromDisk(buffer, count);
    return result;
}
```

In this example, I don't use the BAudioSubscriber to get out of the user data, but it shows how you can get access to the object from within the stream function. RdFromDisk is one of my routines that fills the buffer with data from a sound file. It'll return TRUE if there is more data left in the file, FALSE otherwise. This is important because if the stream function returns TRUE, it'll continue to get called with buffers. If it returns FALSE, then it exits the stream.

Which brings us to the completion function. Whenever a BAudioSubscriber exits the stream, the completion function is called. This provides a chance to do any cleanup that's necessary.

Here's the prototype for the function:

```
long completion_function
        (void * userData, long error);
```

If the object exited the stream normally, perhaps by returning FALSE in the stream function, then the error will be set to NO_ERROR. Otherwise, it'll give you an idea of why the completion function is being called.

Well, that's the quick look at playing sound. There's obviously a lot more to it than I could cover here. If you're really interested in sound on the BeBox, download the source code for BeBop. It's a multichannel sound player for the BeBox that's available at http://www.mojosoft.com.

The R in a Circle, Billie Holiday's Gardenia, and Other Trademark Mysteries

BY CHARLES B. KRAMER

Few symbols of legal arcana are as well recognized — and as frequently misunderstood — as the omnipresent R in a circle. Many who know that ® has something to do with trademarks do not know how it differs from the abbreviation ™ or from the less familiar ℠ or from the © symbol that is often incorrectly placed near logos and slogans.

The more one seeks to divine the trademark mysteries, the more mysterious they may appear. Words may be available for use as a trademark (Rocket Games™), but not available for use as a corporate name (Rocket Games Incorporated) — or vice versa. Federal registration of a trademark takes time and costs money, but it doesn't assure superior rights over a prior user of the mark who did not register it. And yet the symbols ®, ™, ℠, and © float above product names and logos on interactive titles, Web sites, film credits, milk cartons, eyeglass frames, and everything else.

Understanding how these symbols differ, when they should be used, and what aspects of their use are unique to interactive development requires some background.

What is a Trademark?

A trademark is a word, graphic, sound, smell, or other symbol that identifies a source of goods. To identify a source, a mark must remind people of a particular company. The process is the same as the one that makes a particular perfume remind you of someone you know. Similarly, Billie Holiday was known for the gardenia she wore in her hair, and vaudevillian Jimmy Durante was known for punctuating jokes with the assurance "I got a million of 'em!" In all instances, the message to the world is the same: *see this* or *hear this* or *smell this*, and *think of me*.

A mark must also make people think that all similar products using the mark are from — or are authorized by — the same company. Holiday's gardenia, for instance, did not legally meet the standard. Another singer wearing one might remind people of Holiday but would not cause them to think that the other singer either was Holiday or was authorized by Holiday to wear the gardenia.

While trademarks most commonly take the form of a word ("Xerox" for products of Xerox Corp.), they can consist of almost anything capable of identifying a source. A picture (Planter's monocled Mr. Peanut), a sound (the three musical notes associated with Nabisco!), or a 3D shape (the hourglass form of a classic Coca-Cola bottle) can qualify. In some instances, even a color (the pink of Owens-Corning spun glass insulation) or a fragrance (the distinctive scent added to a brand of yarn) can also be a trademark.

This raises the possibility that a computer game character or other interactive object may function as a trademark. A game character consists of tangible elements (the character's name and the shapes and colors of its costume) and intangible interactive elements (keystrokes that cause the character to do a back flip). In general, courts find that tangible elements are protectable as trademarks but intangible ones are not. One court, for instance, considering whether an animated cartoon character infringed the trademark of the *Aquaman* comic book character, stated that:

it is difficult to see how such intangible qualities, having an infinite number of possible visible and audible manifestations, can achieve that . . . consistency of representation that would seem necessary to constitute a [trademark] symbol in the public mind.

This reasoning holds open the possibility, at least, that intangible qualities that don't lend themselves to "infinite manifestations" — like, say, a repeated and distinctive jig by an animated rabbit — can become associated with a particular source and function as a trademark.

Interactive titles often have multiple sources (in the trademark sense) and may contain multiple trademarks as a result. A title may contain, for instance, separate marks for the title's publisher, independent design team, and distributor. Interactive titles may also use trademarks to identify third parties, such as software engine designers, whose work is used under a license.

Customers do not need to know anything about the source a symbol identifies for it to be considered a trademark. Rather, they must believe that different products containing the same mark are from the same source or are authorized, sponsored, or approved by the same source. *Star Wars* operates as a trademark, for instance, because customers assume that a toy sold under that name is somehow connected with the series of films that begins with the film of that name. The customers do not need to know the legal name of the company that owns the mark or what factory manufactured the toy.

Selecting a Mark

The trick for jazz singers, vaudevillians, and prospective trademark owners is to select a symbol that identifies a source as strongly as possible. Selection involves the following considerations:

▶ The less related a symbol is to the product with which it's associated, the greater its potential to be considered a trademark. Kodak, for instance, is a made-up word unrelated to camera film, and Roach Motel combines the descriptive word "roach" with one that barely suggests a product that traps them. Both are strong marks. Merely descriptive phrases, such as "double-speed" for a brand of CD-ROM player, can't act as a trademark except in very special circumstances.

▶ A trademark is unavailable if it's likely to be confused with a mark already in use or for which federal registration is pending. Whether confusion is likely depends on a number of factors, including the degree of similarity between the marks, the products with which they're associated, their customers, and distribution channels. Other factors may be pertinent, such as whether the pre-existing mark is well known. The best way to uncover potentially confusing marks already in use is to conduct a trademark search (see the sidebar "Trademark Searches" on page 90).

▶ Whether a mark is available outside the United States is determined on a nation-by-nation basis. A mark that is available in the U.S. might be unavailable in England or elsewhere. Marks

Charles B. Kramer has practiced corporate, copyright, and trademark law in New York City since 1982 for clients in software development and other information businesses. He can be reached at interactivity@mfi.com.

that will be used internationally should also be checked for possible undesirable meanings in applicable foreign languages. The Chevy Nova, for instance, is believed to have sold poorly initially in South America because "Nova" sounds like Spanish for "doesn't go."

Protection Strategies

Once a symbol has been selected as a trademark, the next job is to keep others from using it. Strategies include:

▶ Doing nothing except using the symbol as a trademark.

▶ Registering the trademark under state law.

▶ Registering the trademark under federal law.

The longer a trademark is used, the more tightly it becomes associated with the company using it, and the more the company may be able to prevent others from using it on similar kinds of products. Protecting a mark by use is furthered by using consistent versions of the mark (in the case of a word, by using the same typeface and spelling) in the same places on the same products. State registration gives exclusive rights to use a mark within the state at best. It helps to establish when use of a mark began and the sorts of products it is used on. Otherwise, though, state registration is of little use for trademarks that are used nationally, such as those on interactive titles.

Federal registration gives rights in addition to those provided by use and state registration. Among other things:

▶ An owner of a federally registered mark is pre-

sumed to be entitled to use it nationwide and can stop others more easily (except those who used it first) from using the mark in ways likely to cause confusion.

▶ A federally registered mark can become invulnerable to certain legal challenges during the sixth year following registration.

▶ Although applications take around 12 months to become registrations (see the sidebar "Federal Trademark Application Time Line," below), within weeks of being filed they appear in federal trademark searches, which discourages others from adopting the mark.

▶ An application to register a mark in a foreign country within six months of the U.S. filing date can in many countries (such as England, France, and Germany) get priority based on the U.S. date. Foreign trademarks are handled by foreign attorneys, typically under the co-ordination of a U.S. trademark attorney.

An application to federally register a trademark can only be made by those who are already using a mark in interstate commerce or who intend to use it that way within (as a practical matter) 18 months. When an application is based on intent to use, registration requires actual use to be shown eventually (see "Federal Trademark Application Time Line," below). Use in interstate commerce requires a regular sale (not a special one for trademark purposes) to someone in a state or country other than the one in which the business using the trademark is based. Some marks not eligible for ordinary federal registration

may be eligible for registration on the federal Supplemental Register. Eligible marks must be in actual use and capable of distinguishing goods. Registration on the Supplemental Register gives fewer benefits than ordinary federal registration.

A federal trademark application looks simple, perhaps deceptively so. Among other things, it requires:

▶ Information about the applicant, including name and address. Special rules apply to applicants who are partnerships or who first used the mark before becoming incorporated.

▶ An identification of the products that the trademark is (or will be) used on, which defines the boundaries of exclusive use that the registration will provide if granted. While applicants can try for a broad identification (interactive entertainment in machine-readable form), something narrower (games for personal computers) may speed the application process or avoid a conflict with pre-existing registrations. Often, part of the trademark registration process is a negotiation with a Trademark Office Examiner, whose job includes keeping identifications accurate and narrow.

▶ Specifying one or more classes within which the identification belongs. Most software, for instance, belongs in Class 9 but game software belongs in Class 28 (which includes toys and board games).

▶ Applications based on use require the date of the first sale in which the mark was used

FEDERAL TRADEMARK APPLICATION TIME LINE

Once a Federal trademark application is filed, completing the registration process usually takes at least 12, and sometimes many more, months. Here's the approximate time line:

▶ **3 months:** The Trademark Office issues a filing receipt, which indicates that the application has become a public record.

▶ **4-6 months:** The Office indicates whether it has any objections to registering the mark. Some objections can be easily overcome (for instance, by changing how your product is described), but others (such as the existence of a pre-existing similar registration) may be impossible to overcome or may require legal arguments that take many months.

▶ **7-10 months:** Your mark is published in the Official Trademark Gazette, which gives a 30-day opportunity (unless extended) for anyone who doesn't want your mark to be registered (because, for instance, they believe your mark is likely to be confused with theirs) to oppose registration.

If your application is not opposed:

▶ A Registration is issued by the Trademark Office if your application is based on actual use.

▶ A Notice of Allowance is issued if your application is based on intent to use. To obtain a registration, you must begin to use the mark and file a statement of use within six months of the date of the Notice of Allowance. Up to five six-month extensions may be obtained if you haven't begun to use the mark yet.

in interstate commerce. Where the date is not known, the earliest month and year "on or before which" such a sale was made can be used. Because proving this date can be critical in trademark disputes, keeping evidence of it — such as a copy of a purchase order — is a good idea even for companies that don't have immediate plans to file a trademark application.

What a Federal Application Costs

The Trademark Office filing fee is $245 per mark, per class. Filing an application to cover the name of an interactive title that will be sold on CD-ROM and made available for multiplayer use on a Web site, for instance, would require payment of $490 for two classes. Applications for logos require special black and white drawings of the logo that cost around $60. Additional filing fees are required for "intent to use" applications when the actual use needed to obtain a registration has begun.

A trademark attorney can help decide whether the marks revealed by a trademark search conflict with an intended mark, answer an examiner's objections after an application is filed, and plan strategy. Whether such assistance is worth the expense is difficult to generalize. Trademark attorneys are a sometimes elitist bunch who feel their expertise is under appreciated. Many sneer even at the arrogance of attorneys who are not trademark specialists who presume to invade their field.

How much trademark attorneys charge varies enormously. Many charge a flat fee (not including Trademark Office filing fees) of around $400 to prepare and file an application for one mark in one class. Some charge their usual hourly rates for every little step that follows (such as to review the Trademark Office's filing receipt for typing errors), which quickly can become expensive. Others charge a flat fee that covers the entire registration process except for certain extraordinary problems.

After a registration is granted, your dealings with the Trademark Office are not over if you continue to use the mark and want to maintain the registration. Additional filings are required during the sixth year following the date of registration, at the 10th anniversary of registration, and at 10-year intervals thereafter. Additional information about federal registrations is available from the Trademark Office (see the sidebar "Additional Information" on page 89).

Business Names and Trademarks

A business name is any name under which a person or group does business other than their real names. Business names include "doing business as" names and corporate names. They are registered on a state or sometimes county level, and at most are checked against a list of other business names — but not against

a list of trademarks. When a name is not available as a business name, often a slight change — like the addition of the word "international" — is enough to make it available. With trademarks, such small changes are a lot less likely to make available a potentially confusing word.

Business names (Microsoft Corporation) are often trademarks as well (Microsoft®), and the two usages sometimes are difficult to distinguish. A name used with "Inc." or with an address, for example, is a reference to the business itself and is a business name use. A name used without those things, especially if in a stylized form on an opening screen or other place where trademarks typically appear, is probably a trademark use. The distinction between trademarks and business names is important partly because business names are not federally registrable as trademarks. Names and marks, however, can cause the same sort of confusion. If some state allows you to incorporate as Microsoft International, Inc., for instance, you will certainly hear from the lawyers of that other Microsoft if you use the corporation to sell software (and probably anything else).

In short, state clearance of a business name does not mean that others can't object to your use of it. A business name should be cleared with a trademark search, even when you don't intend to use it as a trademark.

®, ᔆᴹ, ™, and ©

So what of ®, ᔆᴹ, ™, and ©? The difference between these symbols will make sense now:

▶ ® identifies federally registered marks and may be used only on products within the class and identification of goods in the registration.

▶ ᔆᴹ identifies unregistered service marks. A service mark identifies a source of services (such as an online service at a Web site) just as a trademark identifies a source of goods. Federally registered service marks get to use the ® symbol.

▶ ™ identifies unregistered trademarks. The ™ and ᔆᴹ symbols don't require permission from anyone to use, and neither has legal significance other than as a declaration to the world, "Hey! I consider this my mark." They are useful to distinguish business names (to which ™ and ᔆᴹ should *not* be added) from marks (to which adding one or both, as applicable, is a good idea). They may also be used on marks for which federal registration is pending.

▶ © is part of a copyright notice, which has three parts: © [year of first publication] [name of copyright owner]. Copyright notice is not appropriate for a product name or slogan since neither is protected by copyright. © without a year or name is often placed near a trademark, which looks properly mysterious but is neither trademark notice nor copyright notice nor any other legal thing.▩

This column provides general information and not legal advice, which requires an evaluation of individual circumstances.

Two big search services are Thomson & Thomson and Trademark Research Corp. Cantwell and Paxton Inc. is an excellent smaller company that generally charges less for equivalent services but does not provide all the services of the larger companies.

Searches of pre-existing U.S. marks can cover three areas: federal marks, state marks, and a database of unregistered marks (including many business names) that might reveal prior users with superior rights. The cost depends on which of the three you wish to search, on how quickly you need results, and sometimes by the number of classes included in the search. A search of federal registrations and pending applications that produces results in six business days can cost as little as $90. A full search (federal, state, and unregistered marks) on a one-day rush basis can cost $500 or more. Choices between these extremes are available.

A Thomson & Thomson trademark database is also available on CompuServe for searches of federal, state, and several foreign country marks (Go Trademarkscan). The results are in the form of text and so do not include logos or other nontext marks. The cost is $20 for each group of up to five U.S. marks and $25 for each group of up to five foreign marks. While the cost can rapidly add up to more than that charged by a search service, news that the search has retrieved nothing costs only $1.

So which search is best? The answer depends mostly on how much money you have and how much risk you're willing to take that the mark you want to use is already in use by someone with superior rights.

▶ **Thomson & Thomson** 500 Victory Rd., North Quincy, MA 02171; vox 800.692.8833, 617.479.1600; fax 800.543.1983, 617.786.8273; Internet tthub@tiac.net. **Reader Service #152**

▶ **Trademark Research Corp.** 300 Park Ave. S., 8th Fl., New York, NY 10010; vox 800.872.6275, 212.228.4084; fax 212.228.5090; Internet info@cch-trc.com. **Reader Service #153**

▶ **Cantwell and Paxton Inc.** 2001 Jefferson Davis Hwy., Ste. 1003, Arlington, VA 22202; vox 703.415.1616; fax 703.415.1618. **Reader Service #154** ▩

Bad Mojo

A Bug's Eye View of Macintosh-Based Game Graphics

BY MARK GIAMBRUNO

Every once in a while, a game comes along that pushes the envelope in concept and execution. *Bad Mojo* is one of those games. It's by Drew Pictures (which recently merged with Pulse Entertainment), makers of *Iron Helix*.

One thing that makes *Bad Mojo* unique is that the player assumes the role of a human magically transformed into a cockroach. Players must guide their animated roach through the dangerous wilds of a seedy apartment and bar, solving bug-sized problems and gathering clues about the building's occupants, the mysterious Oracle, and their own unfortunate circumstances.

According to director Vinny Carrella, the idea emerged from a brainstorming session during which company president/producer Drew Huffman told a story from his childhood. He awoke one morning with a terrible pain in his ear and went to the doctor. As it turned out, it was not an ear infection but a baby cockroach that had crawled into his ear canal. Yuck. Reminds you of that old *Night Gallery* episode with the earwig eating through some guy's brain, huh? Anyway, everyone at the meeting seemed to have had some kind of negative experience with members of the six-legged Blattidae family, and the thought of doing a game from this perspective was unusual and interesting. Thus *Bad Mojo* (evil magic) was hatched.

Into the Pit of Despair

So what did your character do to deserve being turned into one of the least popular members of the insect realm? Embezzle scads of cash from the university where you worked as an entomologist, for starters. Then there's the matter of your mother's unusual locket, the one

Mark Giambruno is the creative director of Mechadeus' first two multimedia titles, *Critical Path* and *The Daedalus Encounter*. His current projects include 3D modeling for an upcoming product from [NON-DIS-CLOSURE], a graphic novel about [SE-CRET], and an interesting new project that [HE COULDN'T TELL YOU ABOUT EVEN IF YOU TORTURED HIM WITH PLIERS AND A BLOW TORCH]. Err, just kidding about that last part.

you almost forget in the rush to leave your nosy landlord and squalid apartment behind. Seems this particular piece of jewelry has magical capabilities, and the next thing you know, you've sprouted antennae and find yourself in a weird room with six tunnels leading out of it. Each tunnel leads to a different section of the building, including your apartment, the bar, the landlord's room, and the kitchen.

Movement is controlled by cursor keys, sending your animated roach skittering across one of more than 800 screens. The roach moves very realistically, controlled not only by your decisions but also by "behaviors" like antennae twitching, naturalistic halting movement, and getting hung up on sticky surfaces.

In the game, you are more or less restricted to actions that a real-life roach could perform, and most of the puzzles are obstacles that make sense on an insect-sized scale. These include negotiating a maze of pipes, moving objects to form a bridge, and fending off predators. Some obstacles require creative thinking or influencing the video characters in the game.

During your explorations, you come across other denizens of the dark. These creatures may help or harm you, depending on what they are and how you respond to them. At certain points in the game, these creatures provide animated or video clues about the environment or the other characters. These fill you in on the storyline or provide valuable clues to solving the puzzles.

As you might expect, the object of the game is to regain your human form, but there's more to it than that. The greater object is to regain your *humanity* as well.

Makin' Mojo

Multimedia is definitely the keyword for *Bad Mojo*. The project required multiple disciplines

Figure 1. *Bad Mojo* combines 3D models (like this roach) and photography for a unique and realistic look.
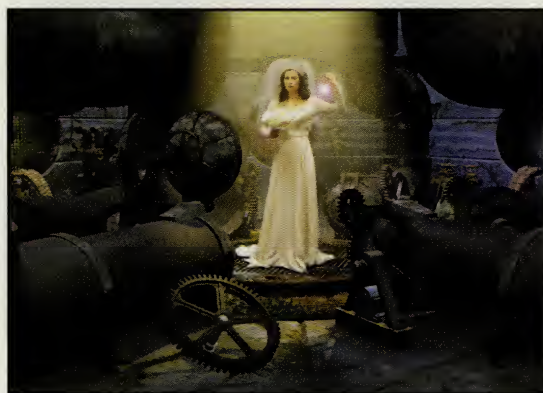

Figure 2. The Oracle provides video and animated clues to guide you through danger.

to create its unique graphic look, including 3D modeling, photography, and video as well as image morphing and retouching.

Since many members of the Pulse Entertainment team are formerly from Macromedia or Paracomp, Macintosh gear and software are the

Figure 3. One of the puzzles involves piloting your roach through a maze of refrigeration pipes in the kitchen.

tools of choice. Most artists there have one or two Power Macs with 32MB RAM for modeling and mapping, plus the company has three 100MB Power Mac rendering stations for the biggest jobs. They also use PCs for coding and demoing.

Most of the 3D work in *Bad Mojo* was done in Macromodel, although auto·des·sys form·Z was also employed, usually when a smooth radius corner or Boolean operations were required.

ElectricImage handled the rendering chores, favored by the team for its speed, ease of use, and deep feature set. "ElectricImage gives us about 95% of the capabilities of an SGI-based solution," Vinny Carrella says. "Even though it's expensive by Mac software standards, it's a bargain compared to those $30,000 workstation software products. It really pays for itself."

One major challenge on the 3D side was that nearly all the images feature extreme close-ups of the environment. The mesh had to be of sufficient resolution that no faceting was discernible in the 640x420 still images without bringing the rendering systems to a halt when producing animations. The variable resolution of spline-based modeling is one reason why Macromodel was chosen. When a close-up was needed, the mesh output resolution was adjusted for smooth curves. For vista shots or animations, a lower res version could be output to ElectricImage without

reworking the model.

Bad Mojo's video is also unusual in that much of it was shot on practical sets constructed inside the Pulse Entertainment building as opposed to using a bluescreen for everything and compositing 3D backgrounds later. Because of this, the modeled versions of the rooms had to match the real sets and vice versa.

One of the most impressive and challenging shots was engineered by technical director Aaron Strain and art director Charles Rose. It's the opening sequence for the game, where a virtual helicopter shot flies us along the San Francisco Bay Bridge (constructed with Macromedia Swivel 3D) and toward a run-down building (form·Z) at its anchorage. The shot continues without pause as the camera flies into the window of the building, revealing video footage of our character packing for his trip. The trick here was getting the movement of the video camera dolly to match the timing and position of the virtual camera in ElectricImage.

The video was shot on Beta SP, digitized using Radius VideoVision Studio hardware and composited with Adobe Premiere and After Effects. QuickTime and QuickTime For Windows were used for the video and animation segments, which were compressed with Cinepak for optimal CD-ROM playback.

Photography was used in many of the close-up screens for realism and to reduce the 3D workload. Objects such as lamps, cigar boxes, and office equipment were brought in by staff members or rented from a prop house to fill out the sets and screens.

Recipe for Roach Gumbo

For a step-by-step look at how a typical environment was created, let's consider the kitchen scene, where your landlord Eddie has been mixing up some catfish gumbo. Delicious.

This area was extensively storyboarded to provide a starting point for the design of each screen. This helped the designers determine what mesh, art, or props each screen would require and how they would be

arranged into a path or puzzle for the user. The layout of the boards created a grid structure that would dictate both virtual camera positions in the 3D models and real camera locations in the shots of sets and props. In many cases, the boards required the 3D models to be squashed or stretched so that they would fit properly into the number of screens allocated.

Technical director Dan Meblin and others used form·Z or Macromodel to create the mesh, sometimes employing both for a given object. For example, the stovetop (see Fig. 6) is composed of numerous Macromodel pieces, while the rest of the stove mesh below it was built with form·Z.

When the mesh for an object was finished, it was output as a DXF file at the required level of resolution. ElectricImage composited the mesh into a completed scene, adding texture maps, lighting, and animation. ElectricImage also has the capability to blend several pieces of mesh into a smooth, seamless whole. This technique was used to assemble the stovetop's 22 pieces into a single object.

Photography was used to add elements like

Figure 5. Vista shot of your landlord's kitchen. A deadly Disneyland for an insect like you.

chunks of food and cooking utensils. Photographer Dann Tarmy set up the props and lighting to match the 3D render using a video camera and live compositing to adjust the elements, then shot on professional tungsten slide film with a 35mm camera. The shots were bracketed (that is, identical shots were taken with the aperture setting notched up and down a few f-stops) to ensure a good exposure. The film was then processed onto PhotoCDs.

The renders and photo elements (see Fig. 7) were adjusted and composited in Photoshop, then went to art director Larry Chandler and artist Johnny Belt for extensive retouching. Most of the game images consist of stills, making it practical to add elements by hand in Photoshop. For example, the grease and food spilled on the stovetop were painted onto the render/photo composite. This approach made it easier to adjust the paths and puzzles as changes occurred during the game design and playtesting. For example, if part of the stovetop grease maze needed

Figure 4. A photo of a dead rat was brought back to life using 'Flo morphing software.

reworking, the changes were quickly made in Photoshop as opposed to altering texture maps, rerendering the affected areas, and then doing any post-production work on the new render.

Some animations involve a combination of 3D and photographic elements. In these cases, portions of the finished screens or their component artwork were reapplied to the mesh as texture maps. In the case of the animation that shows the trapped rat, a metaballs model was made and the photo of the rat used as the texture.

Coding La Cucaracha

The unusual combination of artwork isn't the only unique feature of the game. The C programming team led by Bill Zettler made new inroads as well.

Scribbles is the name of the custom sprite technology developed for *Bad Mojo*. According to assistant producer Alex Louie, "Scribbles are used for the roach, as well as small objects the

---

MAKIN' MOJO: PRODUCT SOURCES

form·Z
auto·des·sys
2011 Riverside Dr.
Columbus, OH 43221
614.488.9777
614.488.0848 FAX
Reader Service #155

Electric Image
ElectricImage
117 East Colorado Blvd., Ste. 300
Pasadena, CA 91105
818.577.1627
818.577.2426 FAX
sales@electricimg.com
Reader Service #156

Macromodel
Macromedia
600 Townsend St.
San Francisco, CA 94103
415.252.2000
800.326.2128
800.288.4797
415.626.0554 FAX
http://www.macromedia.com
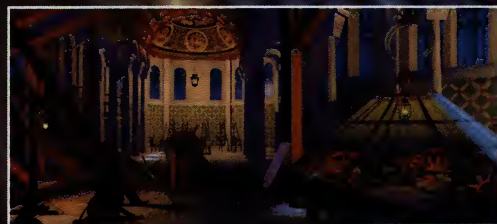Reader Service #157

'Flo (formerly MetaFlo)
The Valis Group
Box 831
Tiburon, CA 94920
800.825.4704
415.435.5404
415.435.9862 FAX
valisgroup@aol.com
Reader Service #158

---

player can move, like cigarette butts or bottle caps."

The behaviors programmed into the roach scribble's movement are among the team's most impressive animations. As you move your roach across the background, it's influenced by the shape and composition of objects it crawls over. For example, if you scale a pipe, the roach seems to rotate as it approaches the edges of the pipe, just as it would in reality. If you attempt to cross some sticky grease, the roach gets hung up and thrashes about wildly.

To create these effects, an unseen grayscale Topography screen is loaded into memory with each background image. This image acts as a realtime bump map, telling the program how to adjust the roach at any given time to match the angle of what it walks over.

There's also a third grayscale screen, the BTC layer (Fig. 8), that triggers events like roach death, animations, and movement behaviors. It also controls the lighting and transparency of the roach and masks it out when it crawls beneath objects. These unique programming effects not only increase the game's reality, but are integral to its play value as well.

Sound design was given careful consideration, breathing life into the scenes with audio effects and original music composed for each of the main stages. A custom WAV mixer developed inhouse for the Windows version allows smooth audio transitions between screens.

All of this comes together with the graphics to create an experience that, as producer Phil Simon puts it, "will make your skin crawl."

Cleaning Up

Bad Mojo is in simultaneous development for both PC and Mac platforms, with other game platforms under consideration. At the time of this writing, Pulse Entertainment is still in negotiation with several publishers, so the shipping date is unclear. With luck, it'll be available in time to infest your home for Christmas '96.
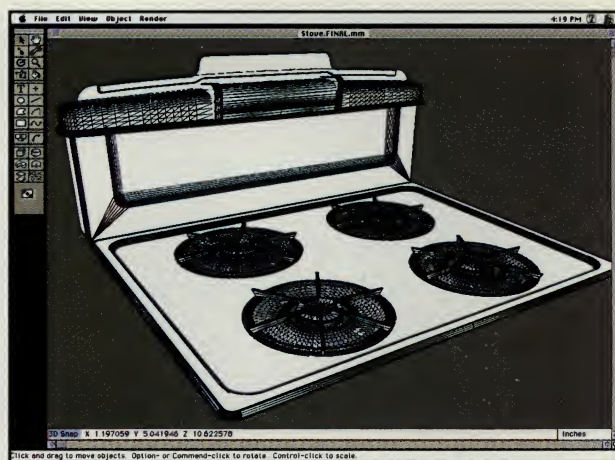
Figure 6. A portion of the stove being constructed with Macromodel.

Figure 7. Render of the kitchen before the photographic elements and retouching are applied.

Figure 8. The BTC layer in action. The roach has a screen effect applied, simulating the dirty glass on the clock, and is masked out where it intersects with the second hand.

Next time, on *Version 5-0*, the latest in retinascalding imagery from the Silicon Graphics crowd. Be there. 🐛

Bad Mojo
Pulse Entertainment (formerly Drew Pictures)
246 1st St., Ste. 402
San Francisco, CA 94105
415.247.7600
415.974.6733 FAX
Reader Service #159

Unleash The Power Of Your Imagination!
ElectricImage™ Animation System

"ElectricImage has allowed us to create intricately detailed environments without the need for high-end Silicon Graphics workstations. The raw horsepower of the program is amazing. No other software even comes close."

Michel Kripalani — President, Presto Studios

I Didn't Know You Could Do That On A Mac!

You don't need an SGI to create dazzling effects. ElectricImage on a Power Mac has been meeting the demands of graphics professionals every day with identical or superior results. You won't get strung out on prolonged rendering times, because ElectricImage is the fastest software-only, per-frame renderer anywhere, beating all other renderers on all platforms — not just on simple tests, but with scenes utilizing 400,000 polygons with transparencies, glow lights, textures, shadows, reflections, motion blur and more! Best of all, ElectricImage is easy to use — you don't have to be a programmer or a Ph.D. to achieve incredible results!

Put the power of your imagination to work with ElectricImage. Call now for a free brochure/flipbook.

Satisfy Your Creative Addiction™

Renderama™ – Network Rendering And Batch Processing

True 3D Inverse Kinematics

Now Up To Twice As Fast

"I-Can't-Believe-It's-Not-Real!" Motion Blur

Anti-Aliasing Levels Of Up To 64x64 Pixels

Easy To Use Inferface

Light Rays – Smoke And Lighting Effects

Smart Deformations

Electric Image
INCORPORATED

Telephone: 818.577.1627
Applelink: ELECTRIC.SLS
Internet: sales@electricimg.com

Images from "Buried In Time: The Journeyman Project 2 CD-ROM Game" © 1995, Presto Studios, Inc.
© 1995 Electric Image, Inc. Character ™&© Electric Image, Inc. ElectricImage, Renderama and Satisfy Your Creative Addiction are trademarks of Electric Image, Inc. Specifications subject to change without notice.

READER SERVICE NO. 21